An Interop Publication

# CONNEXIONS®

## The Interoperability Report

January 1994 Volume 8, No. 1

*ConneXions — The Interoperability Report tracks current and emerging standards and technologies within the computer and communications industry.*

## In this issue:



*ConneXions* is published monthly by Interop Company, a division of ZD Expos, 303 Vintage Park Drive, Foster City, California, 94404–1138, USA.
Phone: +1 (415) 578-6900
Fax: +1 (415) 525-0194
E-mail: `connexions@interop.com`



**ISSN 0894–5926**

### From the Editor

Happy New Year and welcome to Volume 8! This is the eighty-second edition of *ConneXions* since we began publishing in 1987. To date we have put together a total of 2,460 pages. Each and every back issue is still available should you need a particular article. The 1993 index sheet will also be ready soon and will be mailed to you with an upcoming issue. In 1994 you can look forward to more in-depth articles on all aspects of computer networking "from the desktop to the data center."

Just before Christmas we re-located to our new offices in Foster City. Naturally, this involves a few changes to telephone and fax numbers, but since networks are "geographically insensitive" our e-mail address is exactly as before. The parent company, ZD Expos, now houses Interop and Seybold in one location. You can contact us at:

*ConneXions—The Interoperability Report*
303 Vintage Park Drive
Suite 201
Foster City, CA 94404–1138
USA
Phone: +1 415-578-6900 or 1-800-INTEROP
Fax: +1 415-525-0194
E-mail: `connexions@interop.com`

Our first article of the year is an overview of the remote network management (RMON) standards that are part of the *Simple Network Management Protocol* (SNMP) framework. The article is by Jeff Hughes of Hewlett Packard's Network Test Division in Colorado.

Multimedia has become a buzzword in all sectors of the computer industry. In the Internet community, several projects have demonstrated the feasibility of audio and videocasting programs over the existing infrastructure. (See *ConneXions*, Volume 6, No. 6, June 1992). In this issue we look at the *Internet Stream Protocol, Version 2* (ST-II) in two articles. The main application area of ST-II is the real-time transport of digital audio and video packet streams across internets. The first article is by by Ralf Guido Herrtwich and Luca Delgrossi of IBM's European Networking Center and describes the details of the ST-II protocol. The second article, by Chip Elliott and Charles Lynn of BBN Systems and Technologies, outlines a number of ST-II implementations as well as experiences with the protocol and areas for future work.

This month we begin a new series of articles under the heading "Essay." We've asked several individuals to give us their perspective on computer networking. First out are Ed Levinson and Einar Stefferud with two complementary essays on networking "paradigms." We invite all our readers to send us suggestions for future essays.

# The RMON Standards for Network Monitoring

**by Jeff Hughes, Hewlett-Packard Company**

## Introduction

In 1991, the *Internet Engineering Task Force* (IETF) completed work on a new standard that was to be used for remote network monitoring and troubleshooting. This important new standard, called *RMON* (RFC 1271), would soon change the face of *Simple Network Management Protocol* (SNMP)-based network management.

Users were interested in this new standard for several good reasons. First, it was a standards-based solution to their network troubleshooting problems. This meant that buying an RMON compliant device would protect their investment for the foreseeable future. Second, the existence of RMON suggested that the same tools and methodologies could be used to troubleshoot network problems or to monitor networks, regardless of the vendor supplying those tools. Finally, RMON was a solution that allowed users to troubleshoot their networks from a central location. RMON devices were implemented using low-cost probes that could be distributed throughout a user's network.

Other very capable tools exist for the purpose of network troubleshooting. Network analyzers are dispatched tools which are available from several different vendors, including Hewlett-Packard, Network General, Wandel & Goltermann, and others. In contrast to RMON compliant devices, these tools are dispatched to the site of a problem *after* the problem has become apparent to the end user. Another important difference is that there is no standard for the type of measurements that such a device must implement. Thus, different vendors provide different measurements. Although each analyzer is slightly different, they all generally provide excellent diagnostic capabilities, usually with a relatively high price tag.

Despite lacking some of the highly specialized troubleshooting measurements, RMON devices are also used extensively for network problem isolation. RMON-based solutions provide a key advantage because the probes monitor the network constantly, and can provide historical information. Thus, the Network Engineer can determine the behavior of the network before, during, and after a given problem occurs.

Recent developments have made RMON devices even more valuable. The IETF has just finished work on a new Token Ring RMON standard (RFC 1513), which will extend the benefits of distributed troubleshooting to Token Ring networks. This new standard suggests that Network Management Professionals will be able to manage their Token Ring networks using similar tools and techniques to those previously used only on Ethernet networks.

## RMON capabilities

The RMON MIB and the new Token Ring extensions empower the Network Manager to take control of network problems. Regardless of the network technology used, the network manager can obtain statistical data on traffic and error levels occurring on their network. RMON provides statistics in real-time, and also in a historical database that can provide insight into how a problem might have developed. Network Managers tend to use this statistical information as a "wide angle lens" into their network.

For more specific information about network problems, RMON devices provide statistical information on individual hosts (or stations) on the network. This information allows the network manager to "focus in" on specific hosts that might be causing a problem.

RMON devices allow the user to determine the top talkers on a given network, and also the top error generators. Traffic and error statistics may be viewed for a particular host, and also for a particular connection between two hosts on the network.

In order to allow the network manager to troubleshoot at a low level, RMON devices provide packet capture capabilities. The RMON user can create a filter specification to restrict the set of packets he (or she) wishes to capture. The captured packets can then be examined "under a microscope" to gain insight into a specific network problem. Management station software will generally provide packet decoding capabilities that allow the Network Engineer to examine packet exchanges between devices.

Finally, the RMON MIB provides a facility for proactive network monitoring. Network Management Professionals can configure alarms which are to be triggered by specific network events. A typical alarm might be configured to trip when the packets/second level on the network exceeds a predetermined value. Using alarms such as this, the Network Manager can be alerted to a network problem the instant that it occurs.

**Probes**

Due to the amount of processor power required to monitor all packets on the network, most RMON devices today are specifically dedicated to the purpose of network monitoring. These devices are typically called RMON "probes" and are monitored by SNMP-based network management applications. The success of the RMON standards will undoubtedly lead to more widespread implementation in the near future. Expect to see hubs and routers with full RMON capability emerge in the marketplace over the next few years.

**SNMP and RMON**

The RMON MIB and Token Ring MIB extensions were both derived from the world of SNMP-based network management. The model for SNMP-based devices is illustrated in Figure 1.

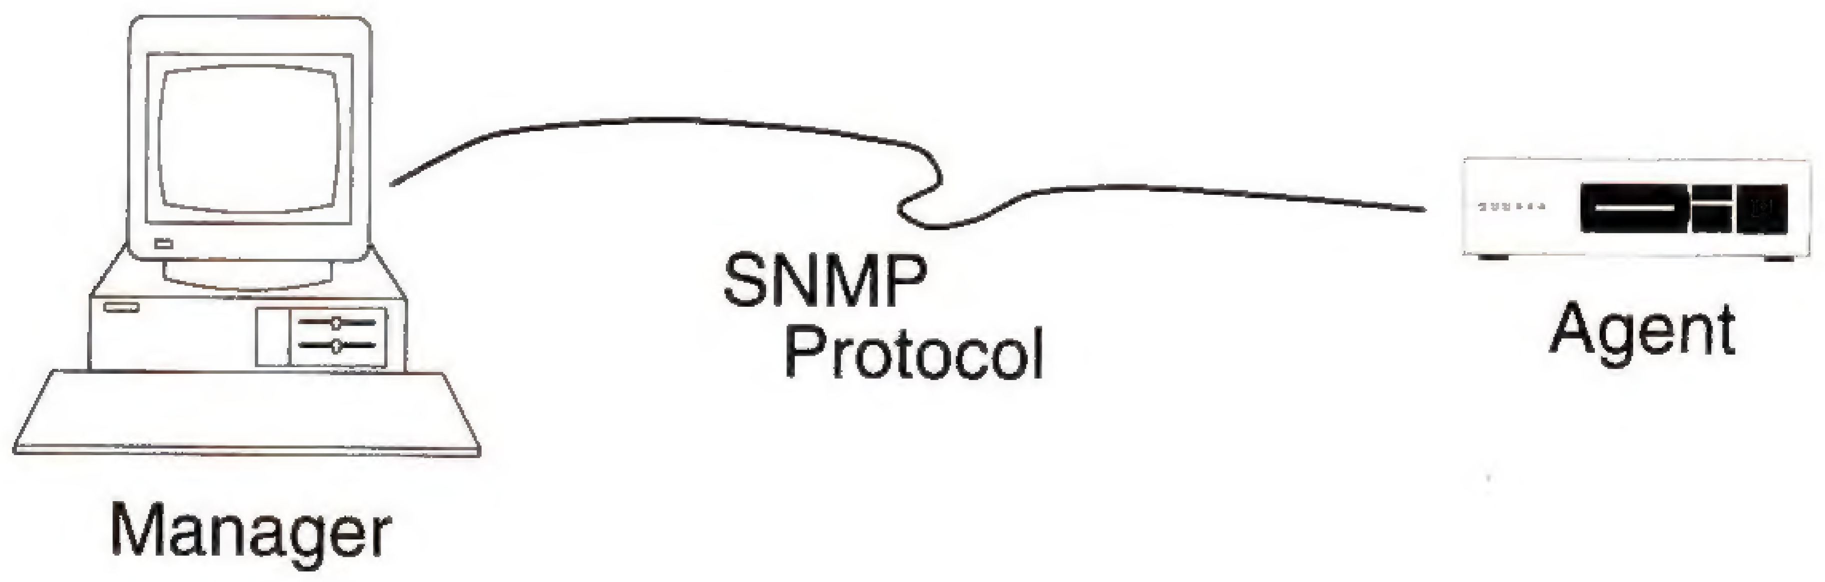


Figure 1: The SNMP Network Management Model.

As an overview, SNMP management is based on a Manager–Agent model, which is basically identical to the more familiar Client–Server model. Using the SNMP network protocol, a management station will "poll" a given agent for data that is maintained by that agent. The agent will respond by sending the data back to the manager, again using the SNMP protocol.

The agent itself is a piece of software running on a network device that is able to respond to SNMP commands. The agent will support one or more MIB structures (MIB is an abbreviation for *Management Information Base*). A MIB is simply a well-defined data structure that may be accessed by other SNMP speaking devices.

*continued on next page*

## RMON Standards for Network Monitoring *(continued)*

Let's take a specific example. A typical Token Ring RMON probe will be a dedicated piece of hardware that attaches directly to your Token Ring network. The probe will monitor all network activity on the ring that it is attached to. As it is monitoring, it will keep counts of various network activities and will store these statistics into a MIB that conforms to the Token Ring RMON standard. The user will monitor the Token Ring probe from an SNMP-based management station. The management station will send SNMP packets to the probe to retrieve information that the user has requested for display.

In general, a Token Ring probe keeps a variety of information, and thus will support multiple MIBs. A diagram of the different MIBs typically supported by a Token Ring probe is shown in Figure 2.

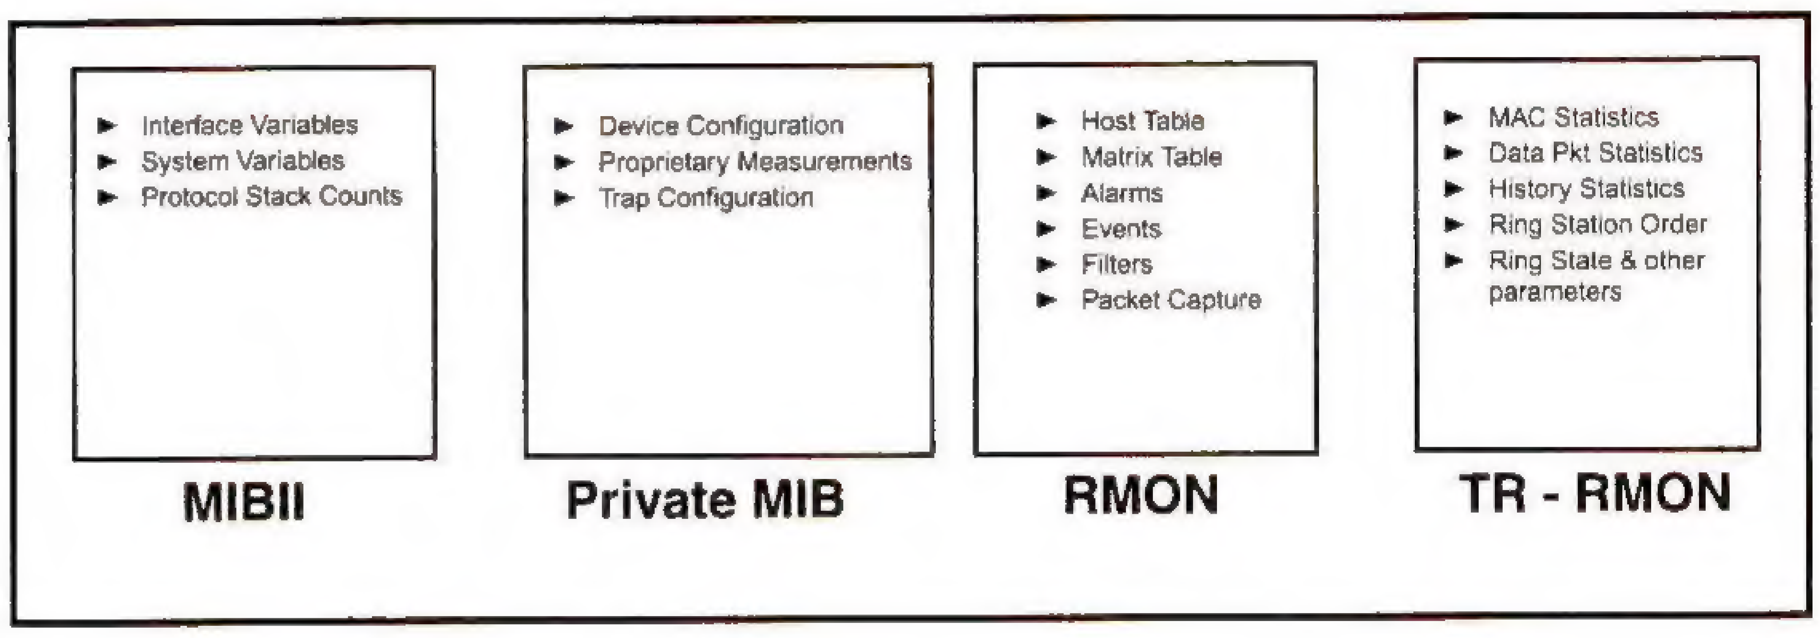


Figure 2: Various MIBs maintained by a Token Ring RMON device.

The RMON MIB is similar to other MIBs in that it is compliant with the SNMP standards. The RMON MIB is also different from other MIBs in some very important ways. First, RMON is a MIB that was specifically designed for network troubleshooting. As such, the data that is maintained pertains to the network itself, *not* the device on which the MIB is located. Other SNMP standard MIBs are concerned with the specific status or throughput of a particular device.

Secondly, the RMON MIB is configurable as to what type of data it should maintain. The user may request that a particular study be performed, or that certain types of packets be captured. The user may specify different intervals for history studies, and may set up very specific alarm conditions.

**RMON specifics**

The original RMON MIB was partitioned into 9 different *groups*. While most RMON probes implement all 9 groups, it is only necessary to implement one group to claim RMON compliance. Figure 3 shows the various RMON groups and a few sample variables from each.

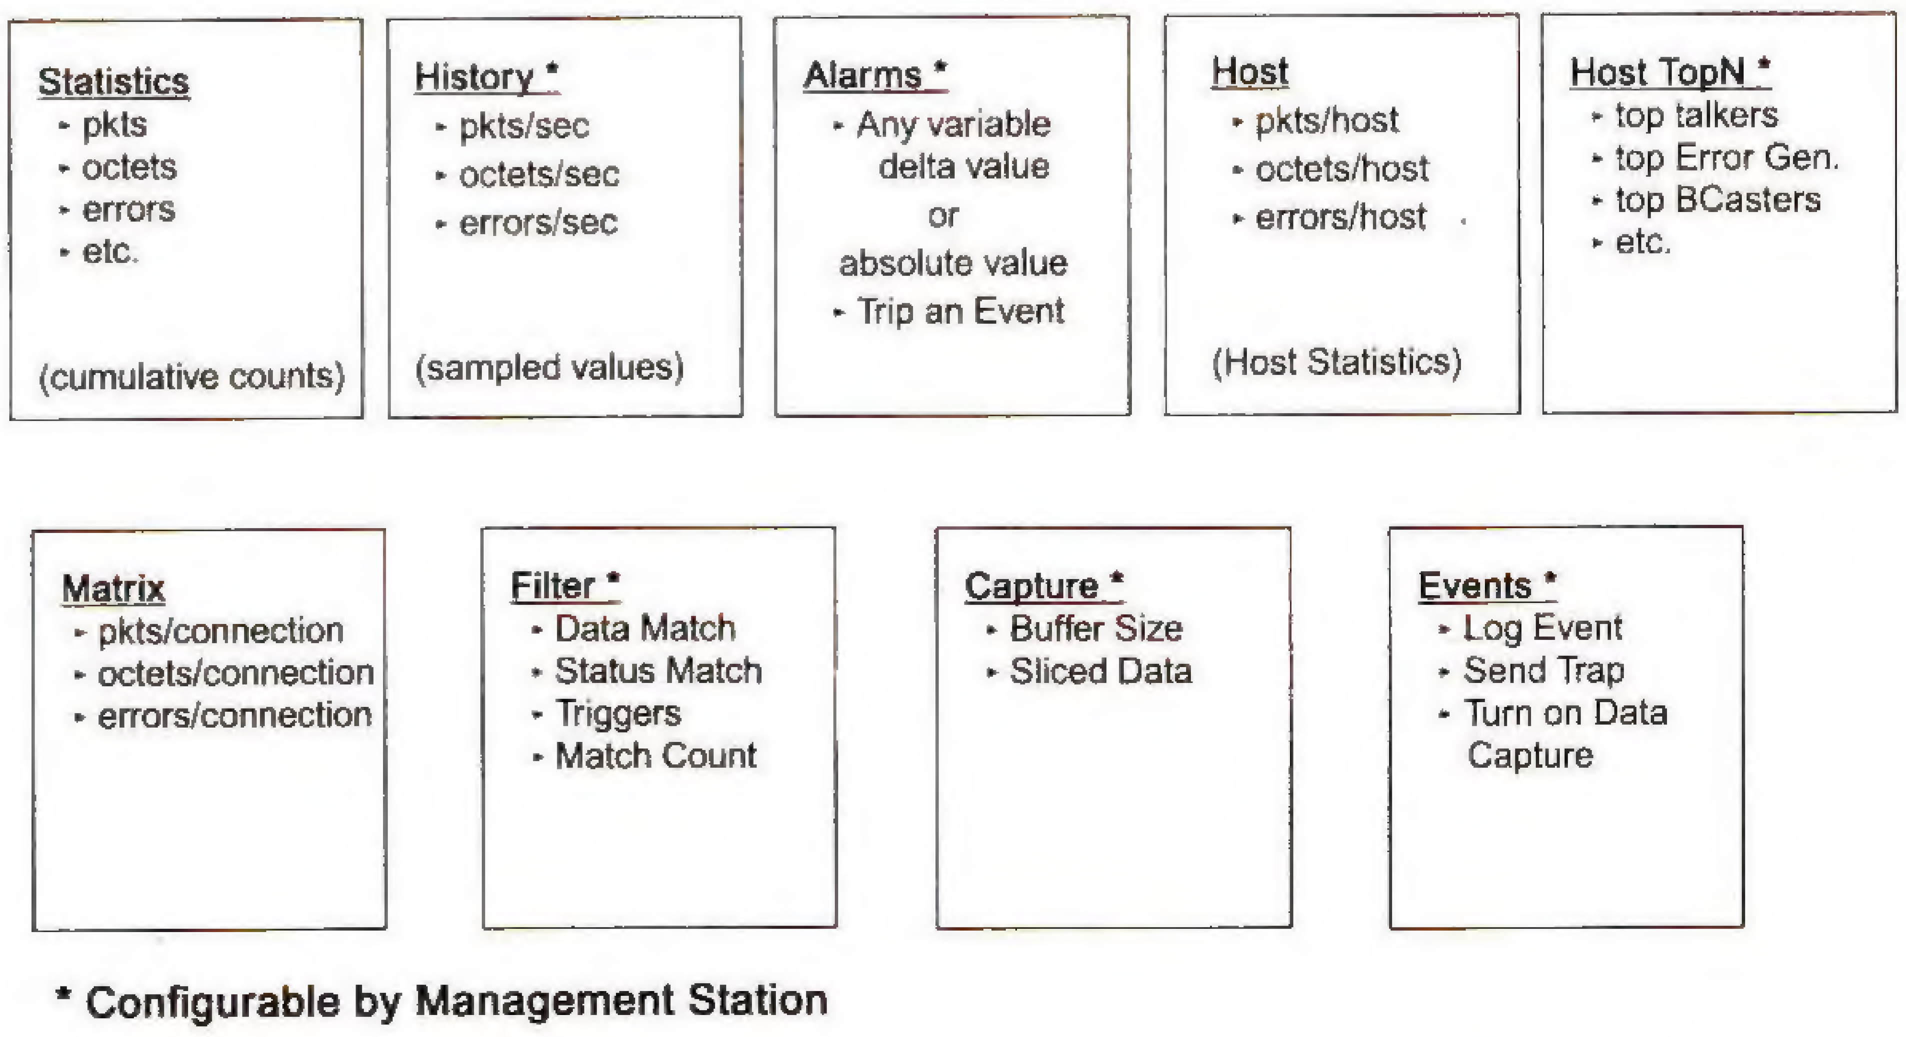


Figure 3: The 9 RMON groups and selected variables from each.

The statistics group contains counts of all network parameters that might be of interest to the network manager. On an Ethernet network, counts of packets, octets, broadcasts, collisions, runts, jabbers and other parameters are maintained. The history group is simply a time based sampling of those parameters in the statistics group. By default, most RMON devices will come with two history studies configured: 5 seconds and 30 seconds. This means that the probe will record the difference in all statistics parameters on each of these two intervals.

The host and matrix groups are also closely related. For each host on the network, the host table maintains statistics such as `inOctets`, `outOctets`, `inErrors`, `outErrors`, `inPackets`, `outPackets`, etc. The matrix table keeps similar statistics on a connection-oriented basis. The `hostTopN` group provides the Network Manager with the ability to ferret out the nodes that generate the most errors, or nodes that generate the most packets. This group effectively adds a sorting capability to the statistics in the host table. The user must generally request a `topN` study, since there are none configured by default.

The filter and capture groups work cooperatively. These groups enable the user to specify a matching criteria for packets that are to be captured and saved in a capture buffer. They keep track of how many packets have been captured, and what to do when the capture buffer fills up.

Alarms are very flexible and may be configured to trip under a variety of circumstances. Statistical counter values, packet matches, and other criteria may be used to trip an alarm. When an alarm trips, it fires the associated event. The event group then determines what action should be taken. A packet can be sent to the management station to notify the network manager of the problem, or the event could simply be saved in the log table on the probe.

## Extending the standard to cover Token Ring

To support Token Ring, several modifications and extensions were needed to the original RMON specification. Basically, four new groups were added, and the statistics and history groups were modified to contain different variables. The Token Ring extensions are illustrated in Figure 4.

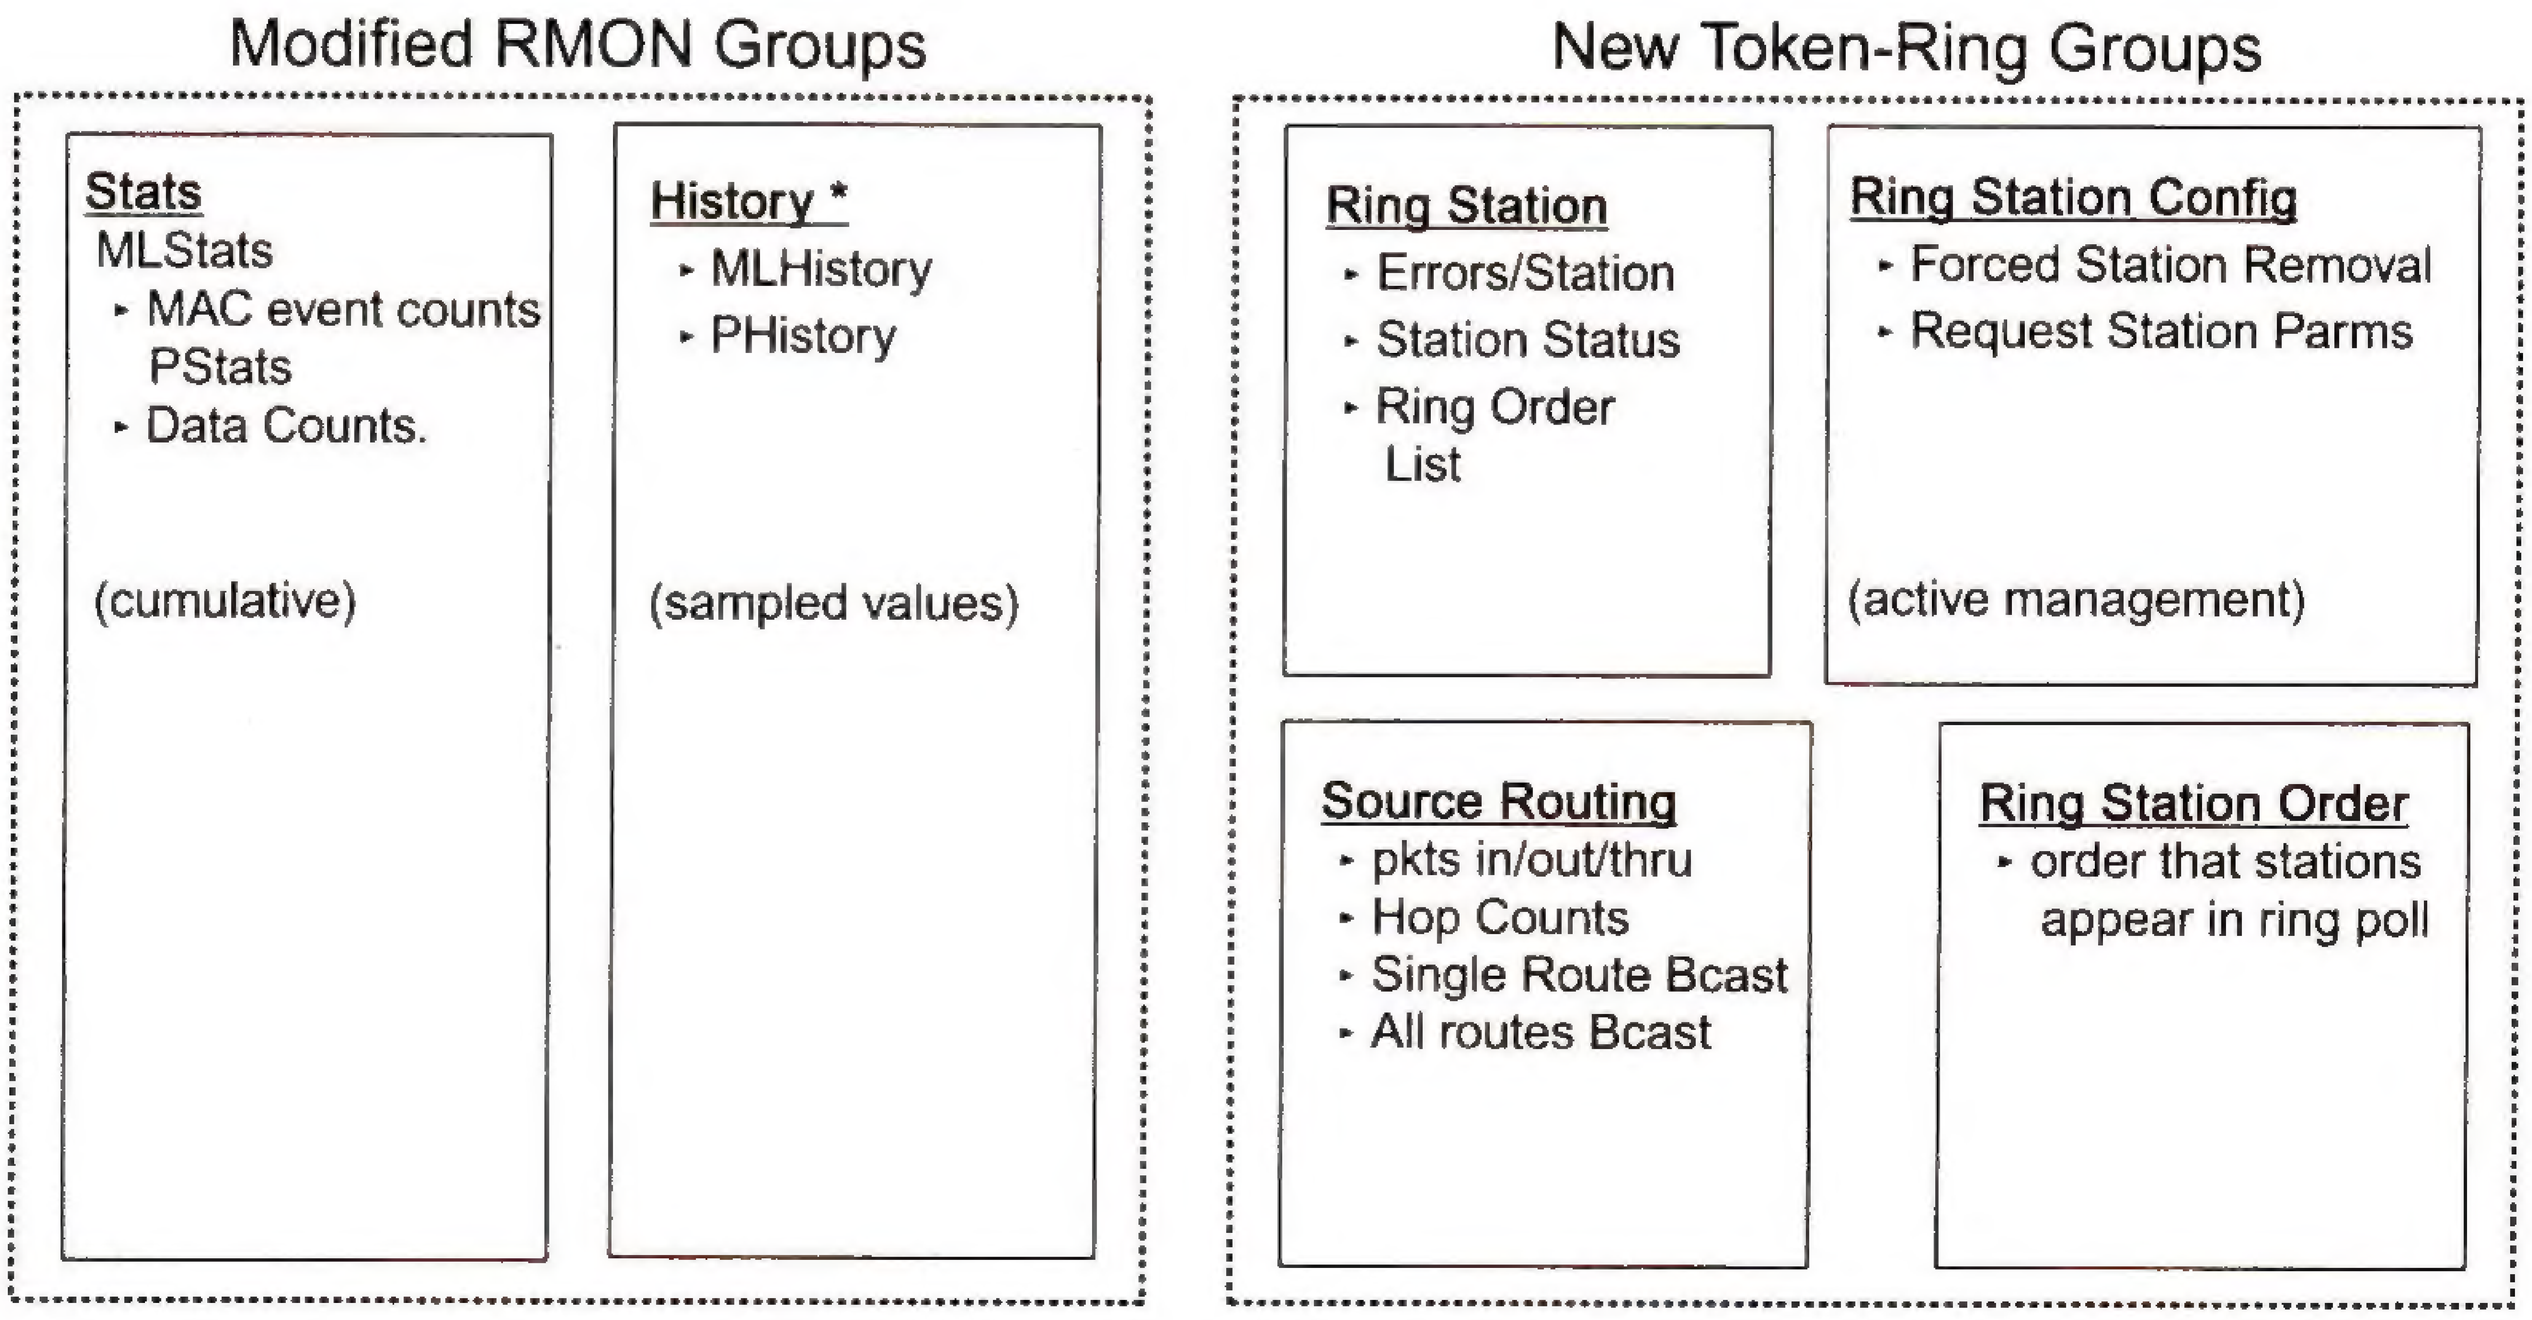


Figure 4: Token Ring Extensions to the RMON MIB.

*continued on next page*

## RMON Standards for Network Monitoring *(continued)*

To start with, the statistics group that was part of the original RMON MIB was replaced with statistics that are more appropriate for Token Ring. Two different flavors of Token Ring statistics are kept: 1) counts related to the packets that are used by the Token Ring protocol (also called "MAC packets") to keep the ring functioning properly; and 2) counts related to the data packets sent between stations that are exchanging information. Again, the history group is simply a time based sampling of the Token Ring statistics parameters.

Four new groups are also defined in RFC 1513. These four groups add significant functionality to the original RMON MIB. The new groups are the *Ring Station Group,* the *Ring Station Order Group,* the *Ring Station Config Group,* and the *Source Routing Group.*

The Ring Station Group essentially keeps track of Token Ring specific information for each station on the local ring. When Token Ring error frames are sent out, this information is recorded against specific stations, to assist with the troubleshooting process. Often times, it is possible to isolate the station or stations causing Token Ring errors by examining the contents of the `RingStationTable.` This table (actually the associated control table) keeps track of the state of the ring, the number of active stations, and the last station known to have sent a beacon frame, among other things.

The Ring Station Order Group provides obvious value in the troubleshooting process. The first thing a network technician needs to know is "What is the order of the stations on the ring?"

The Ring Station Config Group allows for active management of stations on the ring. By setting the appropriate MIB variables, the Network Manager can send a Token Ring frame to a specific station on the ring. In this manner, the Network Manager can "force remove" a station from the ring, or can request that the station respond with its initialization parameters.

Finally, the Source Routing Group contains information related to source routed frames on the network. It keeps track of how many hops a frame had to make on its route from the source ring to the destination ring. It also tracks how many frames came into, went out of, or passed through the ring that the TR-RMON device was attached to.

### Summary

In conclusion, RMON compliant devices are adding significant troubleshooting capability to today's networks. The arrival of the new Token Ring RMON standard will further embellish the arsenal of tools available.

Because RMON is a well-accepted industry standard, Network Managers can feel secure that they are buying equipment that will not become obsolete. Furthermore, they may be able to apply very similar methods when troubleshooting their network problems, on both Token Ring and Ethernet networks.

RMON and Token Ring RMON devices contain a wealth of information that can be used to manage today's highly integrated, multivendor networks.

## References


[1] Waldbusser, S., "Remote network monitoring Management Information Base," RFC 1271, November 1991.

[2] Waldbusser, S., "Token Ring Extensions to the Remote Network Monitoring MIB," RFC 1513, September 1993.

[3] Case, J. D., McCloghrie, K., Rose, M. T. & Waldbusser, S. L., "Introduction to version 2 of the Internet-standard Network Management Framework," RFC 1441, April 1993.

[4] Case, J. D., McCloghrie, K., Rose, M. T. & Waldbusser, S. L., "Structure of Management Information for version 2 of the Simple Network Management Protocol (SNMPv2)," RFC 1442, April 1993.

[5] Case, J. D., McCloghrie, K., Rose, M. T. & Waldbusser, S. L., "Textual Conventions for version 2 of the Simple Network Management Protocol (SNMPv2)," RFC 1443, April 1993.

[6] Galvin, J. M., & McCloghrie, K., "Administrative Model for version 2 of the Simple Network Management Protocol (SNMPv2)," RFC 1445, April 1993.

[7] McCloghrie, K. & Galvin, J. M., "Party MIB for version 2 of the Simple Network Management Protocol (SNMPv2)," RFC 1447, April 1993.

[8] Case, J. D., McCloghrie, K., Rose, M. T. & Waldbusser, S. L., "Protocol Operations for version 2 of the Simple Network Management Protocol (SNMPv2)," RFC 1448, April 1993.

[9] Case, J. D., McCloghrie, K., Rose, M. T. & Waldbusser, S. L., "Management Information Base for version 2 of the Simple Network Management Protocol (SNMPv2)," RFC 1450, April 1993.

[10] Case, J. D., McCloghrie, K., Rose, M. T. & Waldbusser, S. L., "Manager-to-Manager Management Information Base," RFC 1451, April 1993.

[11] *ConneXions,* Two *Special Issues* on Network Management and Network Security, Vol. 3, No. 3, March 1989 and Vol. 4, No. 8, August 1990.

[12] Case, J. D., Davin, J. R., Fedor, M. S., & Schoffstall, M. L., "Network Management and the Design of SNMP," *ConneXions,* Volume 3, No. 3, March, 1989.

[13] Case, J., McCloghrie, K., Rose, M. T., Waldbusser, S., "The Simple Management Protocol and Framework: Managing the Evolution of SNMP," *ConneXions,* Volume 6, No. 10, October 1992.

[14] Rose, M. T., "Network Management is Simple: You just need the 'Right' Framework." In *Integrated Network Management, II,* Iyengar Krishan and Wolfgang Zimmer, editors, pages 9–25, North Holland, April 1991.

[15] Rose, M.T., *The Simple Book: An Introduction to Internet Management,* Second Edition, Prentice-Hall, ISBN 0-13-177254-6, 1993.


**JEFF HUGHES** is a Technical Lead Engineer at Hewlett-Packard's Network Test Division. There he has worked on a variety of network test equipment including dispatched network analyzers, distributed network monitors, and network management consoles. Jeff is also an active member of the Internet Engineering Task Force (IETF), where he contributed to the design of the Token Ring extensions to the RMON MIB. His Internet e-mail address is: `jeff@col.hp.com`

# The ST-II Protocol for Multimedia Communication

by Ralf Guido Herrtwich and Luca Delgrossi,
IBM European Networking Center

## What is ST-II?

The *Internet Stream Protocol, Version 2* (ST-II) is a connection-oriented internetworking protocol that operates at the same layer as connectionless IP. It has been developed to support the efficient delivery of data streams to single or multiple destinations in applications that require guaranteed data throughput and controlled delay characteristics. The main application area of the protocol is the real-time transport of digital audio and video packet streams across internets.

ST-II can be used to reserve bandwidth for multimedia streams across network routes. This reservation, together with appropriate network access and packet scheduling mechanisms in all nodes running the protocol, guarantees a well-defined quality of service to ST-II applications. It ensures that each multimedia packet is delivered in time for use by higher layers in the protocol stack, i.e., within its specified deadline. This facilitates a smooth playout of digital audio and video that is essential for this time-critical data which cannot be typically provided by best-effort IP communication.

Just like IP, ST-II actually consists of two protocols: ST for the data transport and SCMP, the *Stream Control Message Protocol,* for all control functions, used mainly for resource reservation. ST is simple and contains only one PDU that is designed for fast and efficient data forwarding in order to achieve low communication delays. SCMP, however, is quite complex. As with ICMP and IP, SCMP packets are transferred within ST packets as shown in Figure 1.

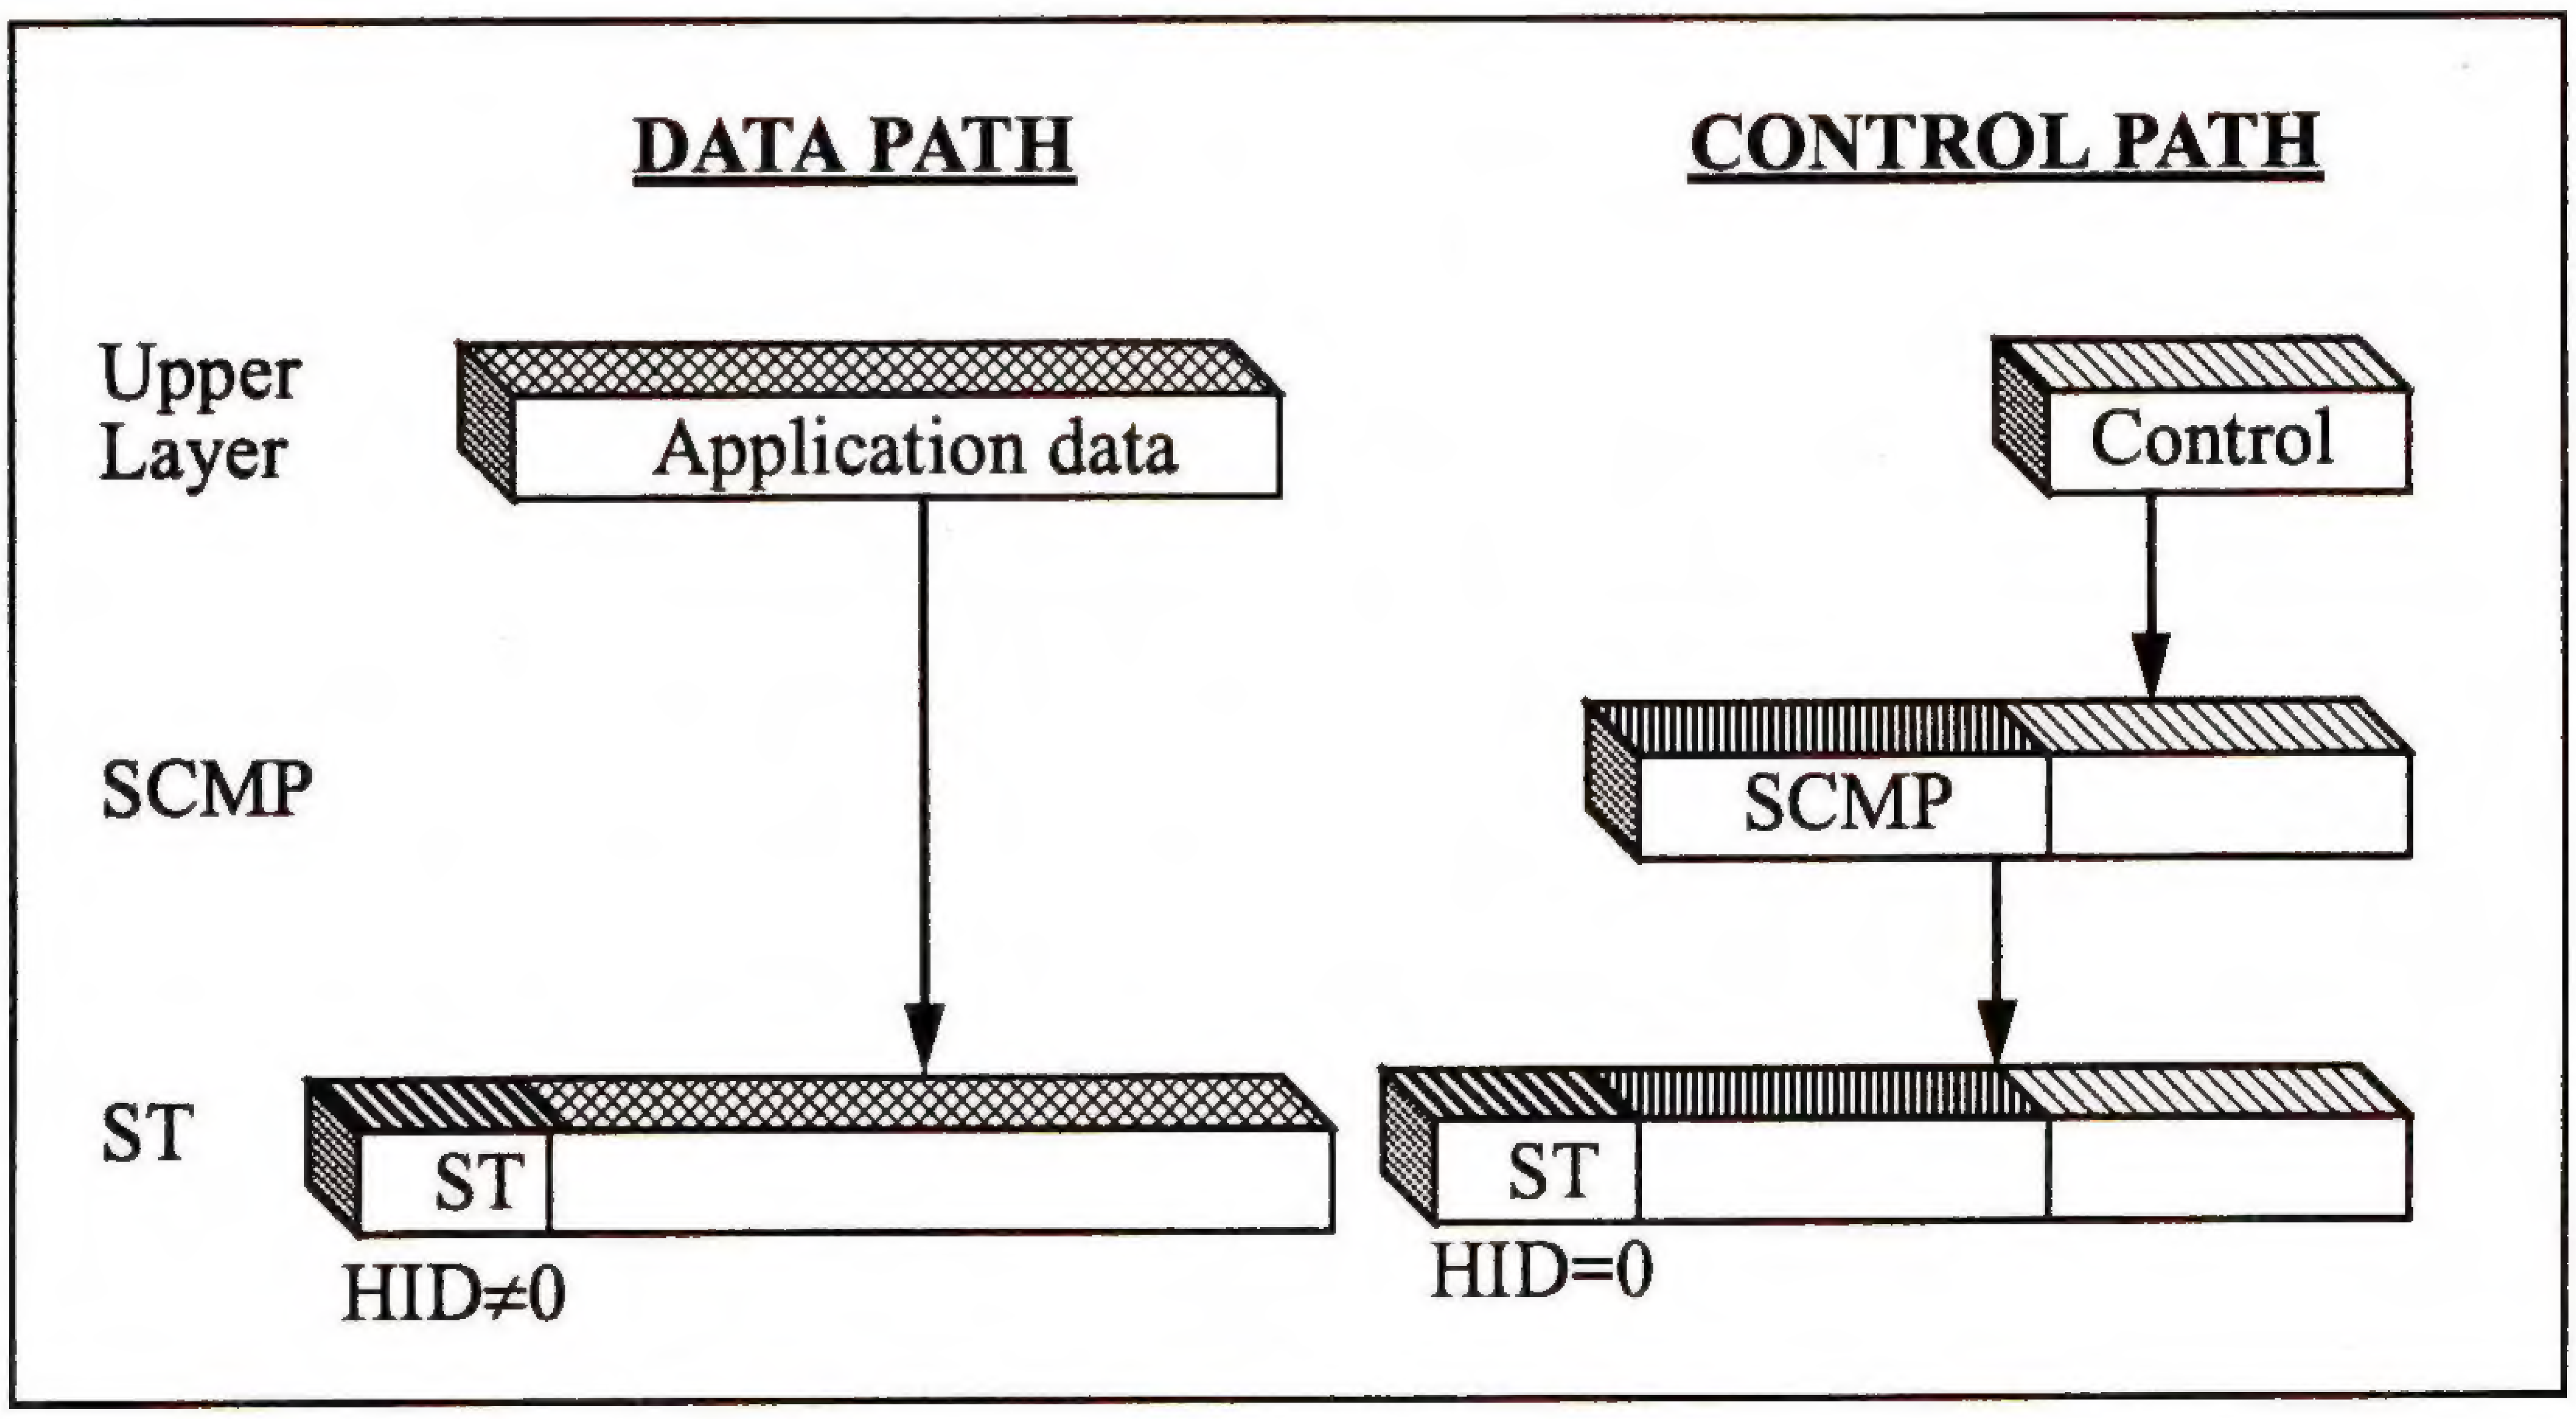


Figure 1: ST-II Data and Control Path

## Protocol history

The first version of ST was published in the late 1970s and used throughout the 1980s for experimental voice and video transmission. The experience gained in these applications led to the development of the revised protocol version ST-II. The revision extends the original protocol to make it more complete and more applicable to emerging multimedia environments. The specification of this protocol version is contained in Internet RFC 1190 which was published in October 1990, classifying the protocol as "experimental."

With more and more developments of commercial distributed multimedia applications underway and with a growing dissatisfaction at the transmission quality for audio and video over IP in the MBONE, interest in ST-II has grown over the last years. Companies such as BBN and IBM have products available incorporating the protocol. The BERKOM project of the German PTT uses ST-II as its core protocol for the provision of multimedia teleservices such as conferencing and mailing. Among others, Digital, HP, IBM, and Siemens-Nixdorf participate in this project. In addition, implementations of ST-II for Sun, Silicon Graphics, Macintosh, NeXT, and PC platforms are available.

Recently, the IETF has started a new working group on ST-II. Its mission is to clean up the current protocol specification to ensure better interoperability between the existing and emerging implementations. It shall also reflect the experiences gained with the current ST-II implementations and applications. The working group is expected to present a revised protocol specification in mid 1994.

**Streams**

Streams form the core concepts of ST-II. They are established between a sending origin and one or more receiving targets in the form of a routing tree. Nodes in the tree represent so-called "ST agents," entities executing the ST-II protocol. Links in the tree are called "hops."

Figure 2 illustrates a stream from an origin to four targets, where the ST agent on one target also functions as a router. Using this Target 2/Router node as an example, we can explain some basic ST-II terminology: The direction of the stream from this node to Target 3 and 4 is called "downstream," and towards the Origin node "upstream." ST agents that are one hop away from a given node are called "previous-hops" in the upstream, and "next-hops" in the downstream direction.

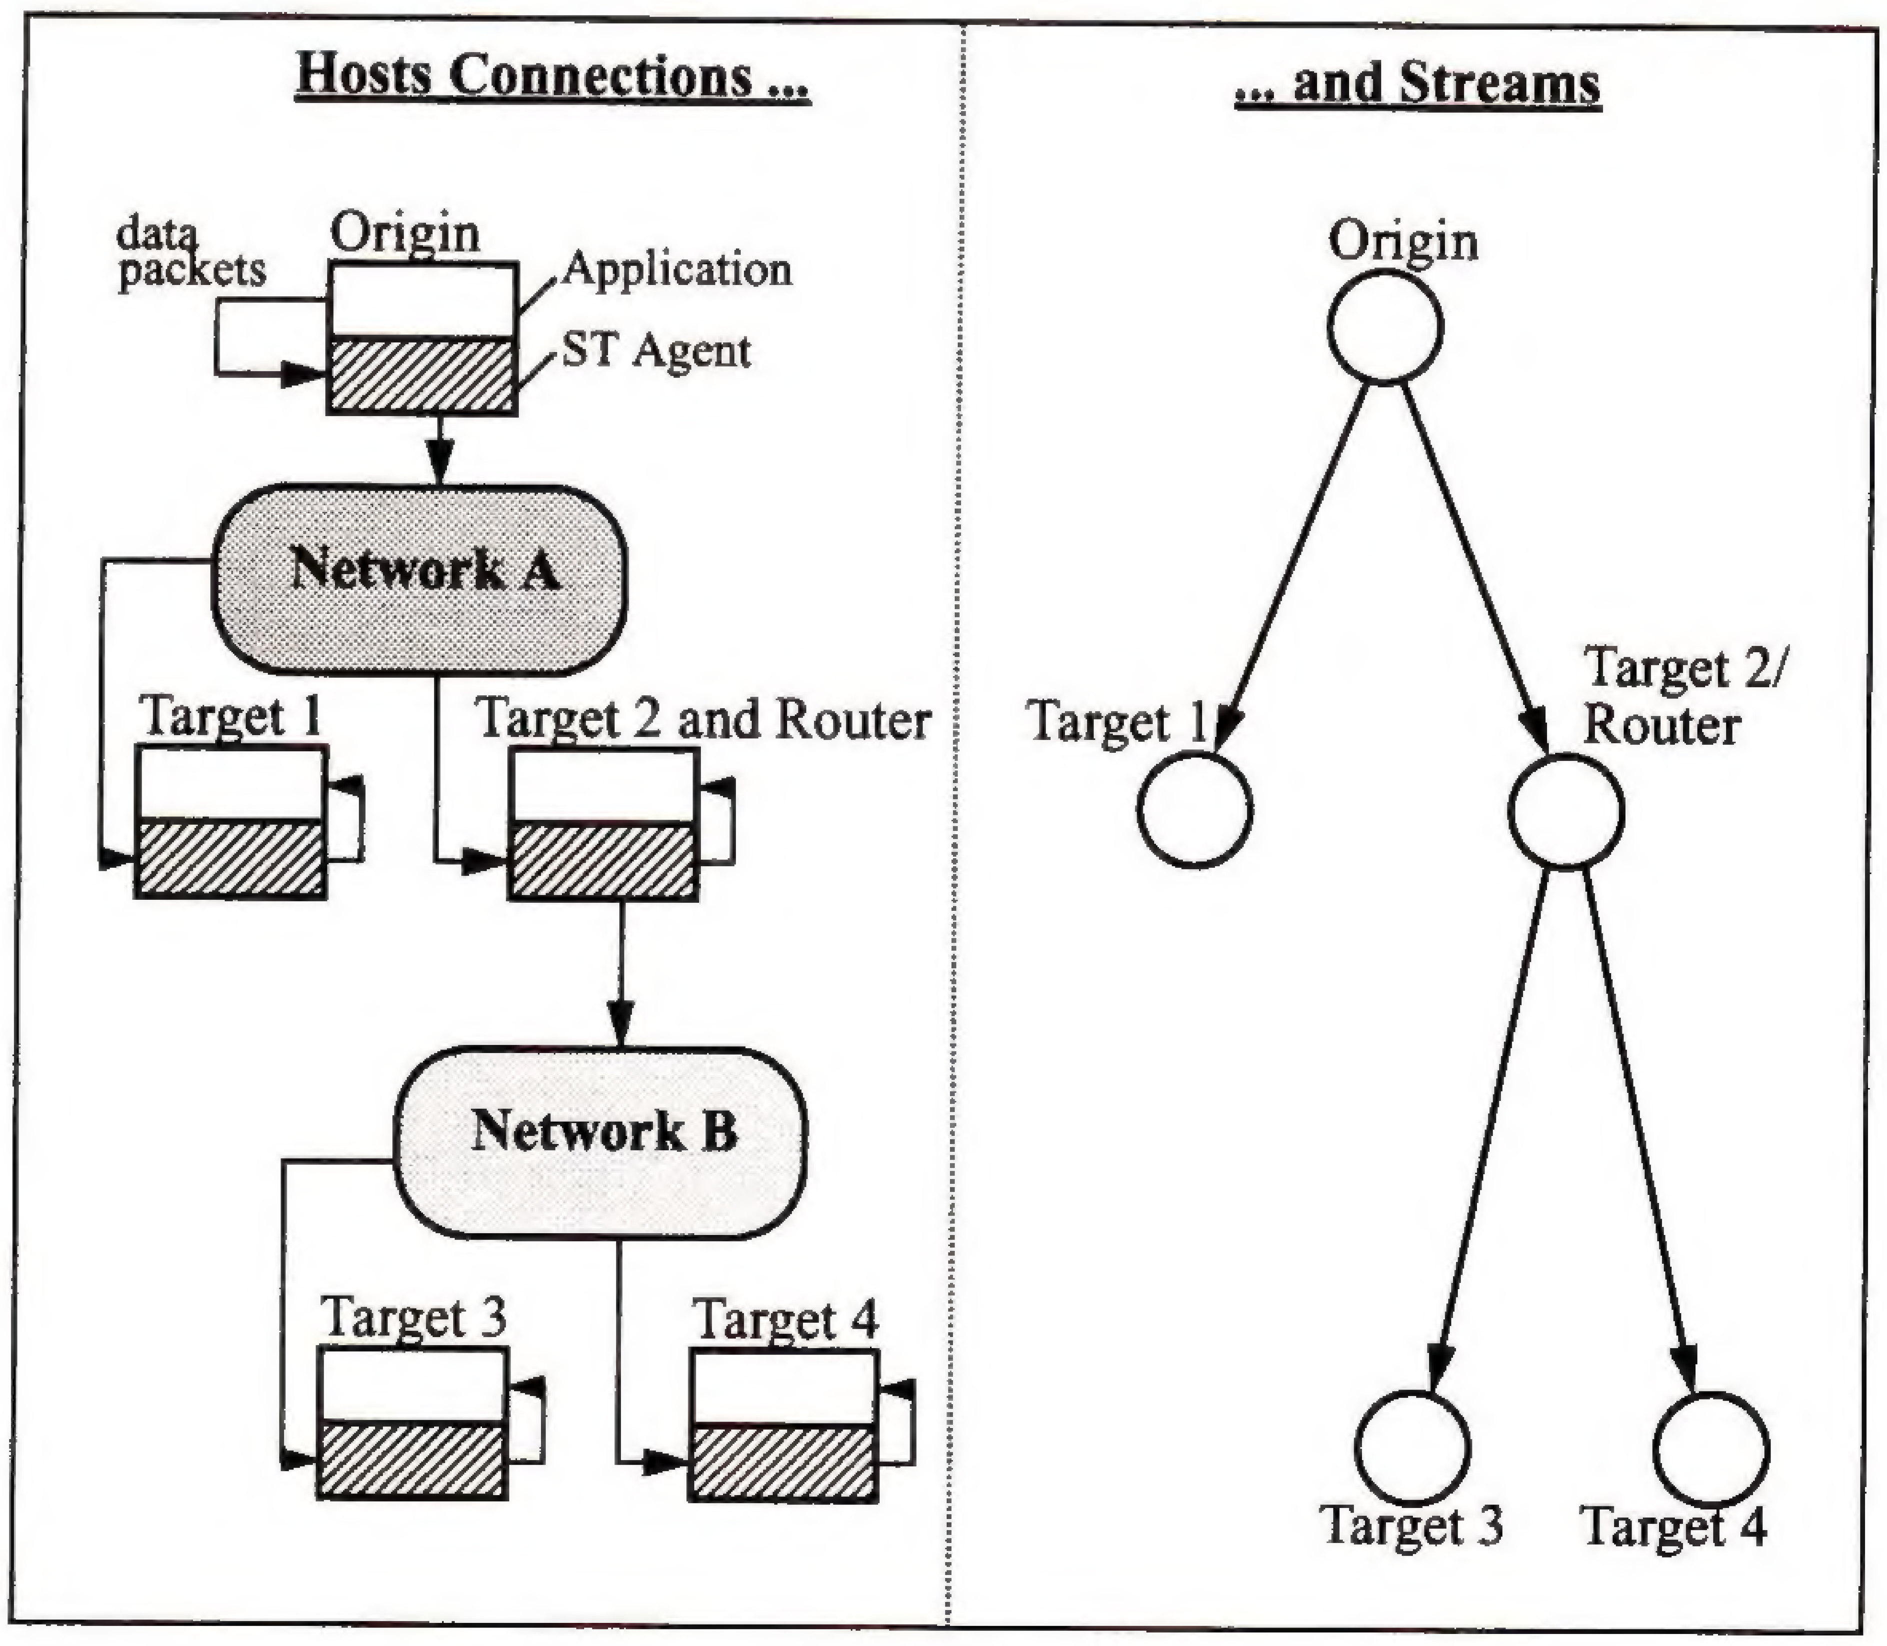


Figure 2: The Stream Concept

Streams are maintained using SCMP messages. Typical SCMP messages include CONNECT/ACCEPT to build a stream, DISCONNECT/REFUSE for the inverse operation, or CHANGE which could be used to modify the stream parameters such as the quality of service or the target set.

*continued on next page*

## ST-II for Multimedia *(continued)*

Each ST agent maintains state information describing the streams flowing through it. It can actively gather and distribute such information. If, for example, an intermediate ST agent fails, the neighbouring agents can recognize this via HELLO messages that are periodically exchanged between ST agents that share streams. STATUS packets can be used to ask other ST agents about a particular stream. These agents then in turn send back a STATUS-RESPONSE message. NOTIFY messages serve to inform ST agents of changes such as a route change.

ST-II offers a wealth of functionalities for stream management such as source routing or route recording. Streams can be grouped together to minimize allocated resources or to process them in the same way in case of failures. During audio conferences, for example, only one person should speak at a given time. Using the group mechanism, resources for only one audio stream from the group need to be reserved. Using the same concept, an entire group of related audio and video streams can be discarded should one of them fail.

**Data transmission**

Normally, data transfer in ST-II is simplex in the downstream direction. Full-duplex communication between the origin and the targets is optional. In full-duplex mode, targets can only communicate with the origin, not with each other.

Data transport through streams is very efficient. ST-II puts only a small header at the start of the user data. This header contains a protocol identification that distinguishes ST-II from IP packets, an ST-II version number, a priority field (specifying a relative importance of streams in cases of conflict), a length counter, a stream identification, and a checksum. These elements form an 8-byte header which can be extended by an optional 8-byte timestamp.

Efficiency is also achieved by avoiding fragmentation and reassembly on router nodes. Negotiations at stream establishment time yield a *Maximum Transmission Unit* (MTU) for data packets on a stream. This MTU is communicated to the upper layers, so that they provide data packets of suitable size to ST-II.

Communication with multiple next-hops can be made even more efficient using MAC Layer multicast. If a subnet supports multicast, a single multicast packet is sufficient to reach all next-hops connected to this subnet. This leads to a significant reduction of the bandwidth requirements of a stream. If multicast is not provided, separate packets need to be sent to each next-hop.

As ST-II relies on reservation, it does not contain error correction mechanisms features for data exchange such as retransmission in TCP. It is assumed that digital audio and video require partially correct delivery only. In many cases, retransmitted packets would arrive too late to meet their real-time delivery requirements. On the other hand, depending on the data encoding and the particular application, a small number of errors in audio and video streams are acceptable. In any case, reliability can be provided by layers on top of ST-II if needed.

**Stream identification**

The identification of streams in ST-II is an issue of interest. In principle, each stream has a globally unique 10-byte identification. It consists of a unique 2-byte number chosen by the origin ST agent, the 4-byte origin IP address, and a 4-byte timestamp. Such a long stream identification could potentially cause long data packet parsing times.

ST agents therefore negotiate an abbreviation for the stream identification when a stream is established. This abbreviation is the 2-byte *Hop Identification* (HID) that is used in every data packet header.

In the case where either a network does not support multicast or multicast is not required, there does not need to be an actual negotiation of an HID—the next-hop simply notifies the previous-hop of the HID to be used. In the multicast case, the sending agent proposes the HID to be used in order to establish a common HID between adjacent ST agents. All receiving agents have to agree on a common HID, otherwise a MAC Layer multicast, where only one physical packet is sent, could not be used. If a proposed HID is already being used by the next-hop, it can propose a set of free HIDs. The sending agent can then choose another HID. This process continues until all next-hops accept the proposed HID. If HIDs are randomly selected there is a high probability that this negotiation terminates within the first three rounds (85.9% after the first, 98.1% after the second, and 99.8% after the third round).

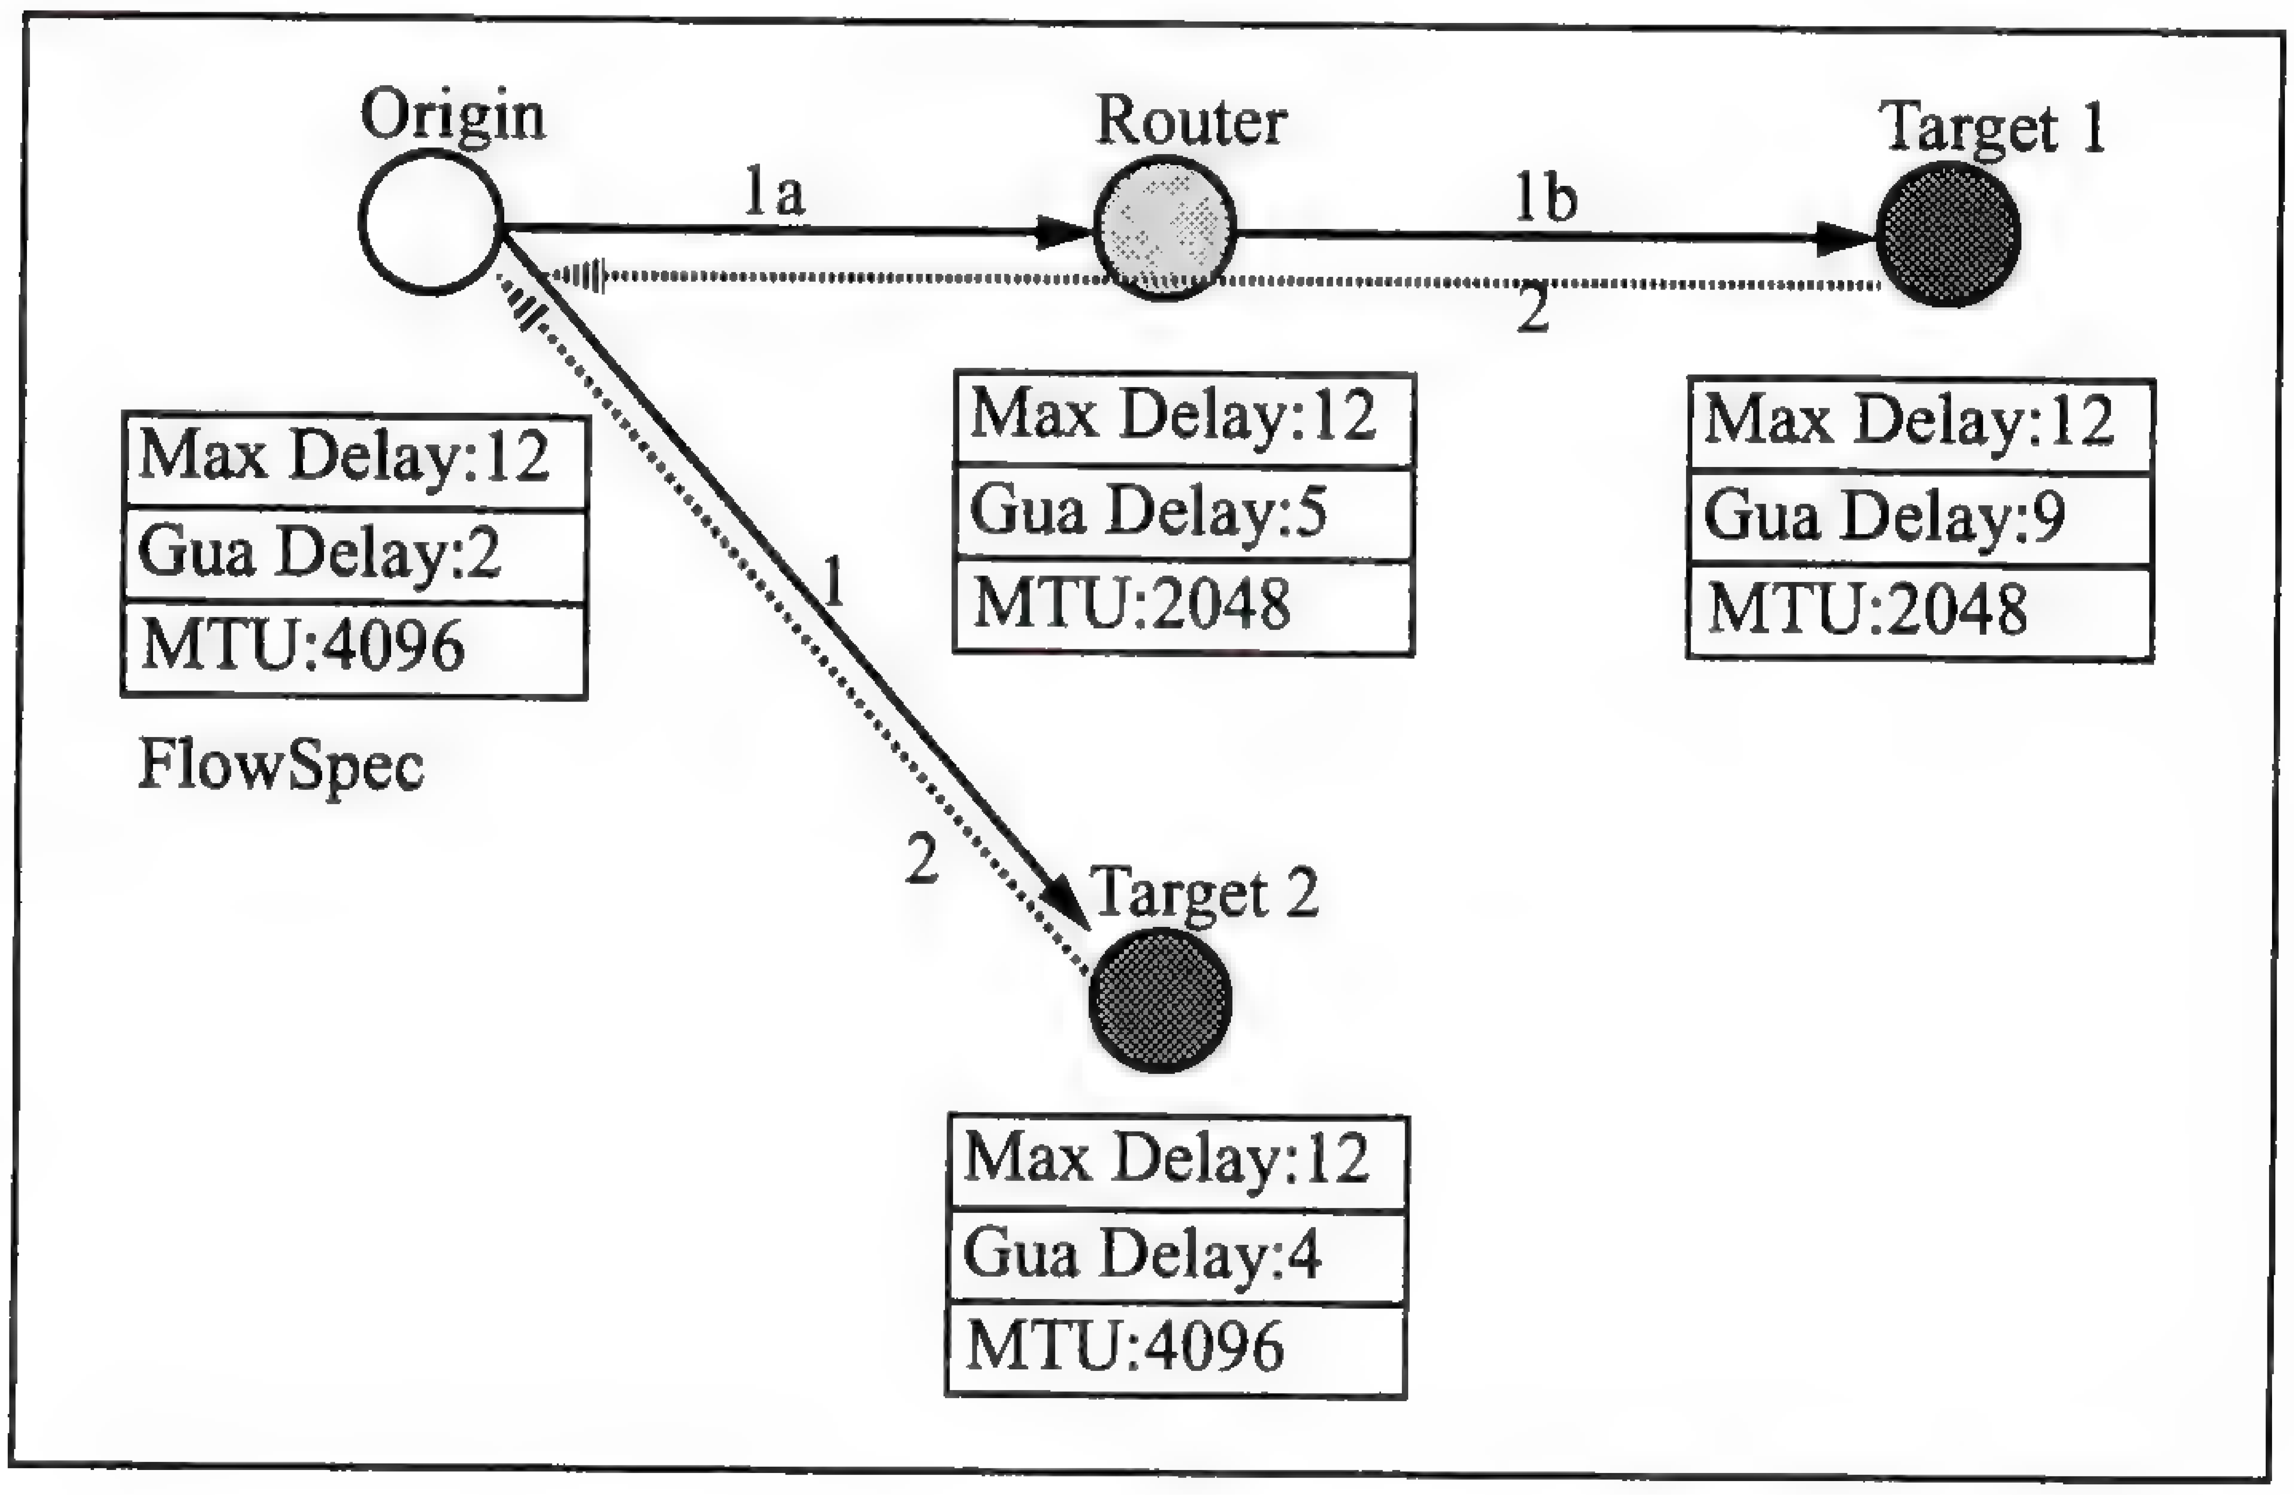


Figure 3: Quality-of-Service Negotiation with FlowSpecs

## Flow Specifications

As part of establishing a connection, SCMP negotiates quality-of-service parameters for a stream. In ST-II terminology, these parameters form a *Flow Specification* (*FlowSpec*, for short) which is associated with the stream. Different versions of FlowSpecs exist and can be distinguished by a version number. Typically, they contain parameters such as average and maximum throughput, end-to-end delay, and delay variance of a stream.

Three kinds of entities participate in the quality-of-service negotiation: application entities on the origin and target sites as the service users, ST agents, and resource managers. The origin application supplies the initial FlowSpec requesting a particular service quality. Each ST agent which obtains the specification as part of a connection establishment message initiates the reservation of local resources by the corresponding resource manager. These resource managers control the usage of CPU capacity for protocol processing, buffer space for storing messages, and bandwidth in the outgoing network. ST-II does not determine how resource managers make reservations and how resources are scheduled according to these reservations, it does, however, assume these mechanisms as its basis.

*continued on next page* 

## ST-II for Multimedia *(continued)*

The FlowSpec negotiation procedure is illustrated in Figure 3 on the previous page. Depending on the success of its local reservations, an ST agent updates the FlowSpec while the connection establishment message passes downstream (for example, keeping track of accumulated delay). The final FlowSpec is communicated to the target application as a basis for its accept/reject decision; the target may finally also modify the FlowSpec according to its needs. If a target accepts the connection, the (possibly modified) FlowSpec is propagated back to the origin which can then calculate an overall service quality for all targets. If all targets in a particular ST-II connection need to adhere to the same FlowSpec, the origin may—during a second phase of connection establishment—issue a CHANGE request to adjust reservations.

**ST-II and IP**

ST-II is designed to coexist with IP on each node. A typical distributed multimedia application would use both protocols: IP for the transfer of traditional data and control information, and ST-II for the transfer of digital audio and video. Whereas IP typically will be accessed from TCP or UDP, ST-II will have new multimedia end-to-end protocols on top of it.

Both ST-II and IP apply the same addressing schemes to identify different hosts and use ARP for address resolution. ST-II can easily be modified to include the longer host addresses of the next generation IP.

ST-II uses the same Layer 2 SAPs as IP. ST-II and IP packets differ in the first four bits, containing the internetwork protocol version number: number 5 is reserved for ST-II (IP itself has version number 4). An ST agent receives a packet over the IP SAP using the first 4 bits of the frame to select ST-II packets.

As a special function, ST-II messages can be encapsulated in IP packets. This allows them to pass through routers which do not run ST-II. Resource management is typically not available for these IP route segments. IP encapsulation is, therefore, suggested only for portions of the network which do not constitute a system bottleneck.

**ST-II extensions**

Several extensions have been proposed and implemented for the ST-II protocol. These include the use of new FlowSpecs and negotiation schemes. Whereas the original ST-II specification only allows for sender-oriented connection establishment, there exist ST-II implementations that also allow targets to connect themselves to an existing stream. Further extensions include concepts such as advance reservation and filtering to provide heterogeneous targets with different levels of service quality within one single stream. It may well be that some of these extensions find their way into the revised ST-II specification that is currently being worked on.

**References**

[1] C. Topolcic, S. Casner, C. Lynn, P. Park, K. Schroder, "Experimental Internet Stream Protocol, Version 2 (ST-II)," RFC 1190, October 1990.

[2] J. Forgie, "ST, A Proposed Internet Stream Protocol," IEN 119, September 1979.

[3] D. Clark, S. Shenker, L. Zhang, "Supporting Real-Time Applications in an Integrated Services Packet Network: Architecture and Mechanism," Proceedings of ACM SIGCOMM, 1992.

[4] L. Zhang et al, "RSVP: A New Resource ReSerVation Protocol," Preliminary Draft, March 1993.

[5] D. Ferrari et al, "Network Support for Multimedia: A Discussion of the Tenet Approach," Technical Report TR-92-072, Computer Science Division, University of California at Berkeley, November 1992.

[6] C. Partridge and S. Pink, "An Implementation of the Revised Internet Stream Protocol (ST-2)," *Internetworking: Research and Experience,* Volume 3, pp. 27–54.

[7] L. Delgrossi, R. G. Herrtwich and F. O. Hoffmann, "An Implementation of ST-II for the Heidelberg Transport System," IBM European Networking Center.

[8] F. O. Hoffmann and L. Delgrossi, "ST-II for the Heidelberg Transport Systems: Protocol, Design, and Implementation," IBM European Networking Center.

[9] R. G. Herrtwich and L. Delgrossi, "Beyond ST-II: Fulfilling the Requirements of Multimedia Communication," IBM European Networking Center.

[10] M. Steenstrup, "Fair Share for Resource Allocation," Draft, December 1992.

[11] L. Zhang, "VirtualClock: A New Traffic Control Algorithm for Packet-Switched Networks," *ACM Transactions on Computer Systems,* Volume 9, No. 2. May 1991.

[12] B. A. Denny, "A Hybrid Algorithm for Combining Best-Effort and Resource Reservation Service," Information and Telecommunications Sciences Center, Report ITAD-8600-TR-93-168, SRI International, June 1993.

[13] J. Gahm, R. R. Hain, and P. Park, "ST-II on T/20 Platform," BBN internal document.

[14] C. Elliott, "High-Quality Multimedia Conferencing Through a Long-Haul Packet Network," forthcoming in Proceedings of ACM Multimedia 93.

[15] C. Partridge, "A Proposed Flow Specification," RFC 1363, September 1992.

[16] C. Elliot and C. Lynn, "ST-II Implementations," *ConneXions,* Volume 8, No. 1, January 1994.

[17] D. Crocker, "The ROAD to a New IP," *ConneXions,* Volume 6, No. 11, November 1992.

**RALF GUIDO HERRTWICH** leads the Distributed Multimedia Solutions Department of the IBM European Networking Center in Heidelberg, Germany. He holds a diploma and doctorate in computer science from the Technical University of Berlin. He has been involved in multimedia research and development since a sabbatical at the International Computer Science Institute at UC Berkeley in 1989/90. In his past three years with IBM, he has worked on making a $15,000 RS/6000 compete with a $150 CD player and a $1,000 TV set. Ralf is one of the three main editors of the ACM/Springer-Verlag journal *Multimedia Systems.* He can be reached as: `rgh@dhdibmip.bitnet`.

**LUCA DELGROSSI** has received his diploma in computer science from the University of Milan in 1986. After work for Olivetti and UC Berkeley, he now is team leader of the Communications and Platform Group at the Distributed Multimedia Solutions Department of the IBM European Networking Center in Heidelberg, Germany. He is one of the principal designers of the Heidelberg Transport System (HeiTS) for real-time multimedia communication and has supervised the multimedia transport system development in the German BERKOM project—both using the ST-II protocol. He currently co-chairs the new IETF ST-II working group. Luca can be reached as: `luca@ibmpa.awdpa.ibm.com`.

# ST-II Implementations

by
Chip Elliott and Charles Lynn, BBN Systems and Technologies

## Introduction

ST-II is an experimental Internet protocol designed in 1990 as a way to multicast "streams" of packets with *Quality of Service* (QoS) guarantees such as bandwidth reservation, delivery delay, and error bounds [1]. It was based on ten years' experience with an earlier protocol, ST [2], that was designed for transmitting realtime audio streams.

Although designed for a range of possible uses, recent ST-II work has concentrated on multimedia transport for such activities as distributed simulations, long-distance training, and in particular video conferencing. Here the bandwidth reservation features of ST-II are used to ensure that audio and/or video packets are not dropped or delayed due to contention from other traffic, as can be the case with standard IP traffic. The Internet multicasts of recent IETF meetings illustrate how badly audio and video transmission can be disrupted when packets are dropped, and the potential benefits of a bandwidth reservation mechanism.

Strictly speaking, ST-II contains two separable mechanisms. First, it allows a sender of packets to establish a multicast "connection" through one or more networks to its targets. New targets can be added at any time, and existing targets can be deleted, so ST-II allows delivery of a media stream to an audience that shifts over time. These connections reserve network resources (such as link bandwidth) along their paths. Second, ST-II allows efficient delivery of data packets down this reserved multicast tree. Thus, when used atop a LAN or WAN with an underlying mechanism that *enforces* the QoS reservations that ST-II establishes, ST-II can transport data with a guaranteed QoS. (Without an enforcement mechanism, of course, reservations cannot truly provide service guarantees. Enforcement is not a part of ST-II proper; rather ST-II reservations are assumed to take advantage of whatever underlying enforcement mechanisms exist.)

## Motivation

This is an opportune time to take a look at ST-II, for several distinct reasons:

- The need for local or long-haul distribution of continuous media streams (e.g., audio and video) is becoming much more pressing as commercial use of multimedia increases.
- Interest is growing in QoS mechanisms. For example, an alternative approach to Internet multicast data transmission with QoS guarantees is gradually taking shape under the aegis of Clark, Deering, Estrin, Zhang, and others [3,4]. This new scheme differs in many important ways from the ST-II approach. In particular, its resource reservation scheme (RSVP) can be characterized as a "receiver based" rather than "sender based" scheme. The Tenet group at Berkeley [5] has tackled the same QoS problems, with notable success, outside of the Internet framework; their work is closer in spirit to the ST-II approach.
- And most important, we now have significant experience with several ST-II implementations and their abilities to transmit multimedia streams.

In fact, we believe that a subset of ST-II is now sufficiently established that it can be used for practical work in distributing multimedia data streams. However, its specification will need some revision before interoperability is possible.

## Existing implementations

Four ST-II implementations are operational. Each is a totally separate implementation and the implementation styles have varied considerably.

• Craig Partridge and Steve Pink have written a public domain ST-II at SICS for the Sun UNIX kernel and documented the experience in a very interesting article [6]. Their implementation adds 6500 lines of **C** to the UNIX kernel and provides a socket interface to user code. Applications are given access to the entire SCMP (control) and data packets, including headers, which allows maximum flexibility in application design, although at the cost of requiring the application to correctly manage all peer-to-peer interactions. This implementation currently supports Ethernet traffic (though without any resource management mechanism) and Fore Systems ATM links.

• Luca Delgrossi, Ralf Herrtwich, Frank Hoffmann and others at IBM in Heidelberg, Germany have created an ST-II that runs entirely within user space on OS/2 and AIX platforms. Here ST-II is used as the underpinnings of a full, production-quality multimedia system. Their version runs over a Token Ring LAN and FDDI with a centralized mechanism for resource (bandwidth and delay) management; above ST-II they run the HeiTP multimedia transport protocol. Their implementation and experience is documented in an excellent set of papers [7, 8, 9].

• Charlie Lynn and Ken Schroder of BBN created another public domain version of ST-II for Sun platforms. This version did not require kernel modifications, although network driver code was modified to provide the resource management functionality assumed by ST-II. The socket-based API differs somewhat in philosophy from the Partridge and Pink implementation in that the user interface hides actual packet formats and the distinction between control and data packets. This implementation is currently installed on workstations on the DARTNET research internet, supports Ethernet and T1 interfaces, and has been used to carry wide-area audio and video traffic. In addition, it has served as a testing tool for *Fair Share* [10], *VirtualClock* [11], and *SFQ* [12] resource management mechanisms.

• A second BBN team led by Josh Gahm has designed an ST-II implementation for the T/20 Router [13]. This implementation is currently being phased into production use on the *Defense Simulation Internet* (DSI). It employs the bandwidth reservation and multicast services provided by its underlying WPS packet switches and deserves note as the only ST-II implementation now being deployed in an internet router on a worldwide scale. At present, a distributed videoconferencing application runs atop ST-II on the DSI [14]; distributed simulation software and security streams are due to follow.

## Implementations in the works

At least nine more implementations are underway and should be available by the end of the year. Most are European efforts related to the BERKOM project, which aims to provide a multimedia infrastructure atop B-ISDN, and most (but not all) take a public domain implementation as the starting point.

The first five are designed for practical use and motivated by immediate need for a multicast protocol for continuous media streams that provides QoS (bandwidth and delay) reservation:

*continued on next page*

## ST-II Implementations *(continued)*

• *Wellfleet:* Paul Goransson has completed a prototype implementation for the FN, LN, and CN routers, and is now designing the production implementation. Wellfleet's ST-II is designed to work on the Defense Simulation Internet and hence interoperates with the BBN T/20 Router version. It supports two resource enforcement mechanisms, VirtualClock and a simpler scheme based on high and low priority queues, over both Ethernet and Frame Relay circuits. There is some interest in supporting ST-II over ATM.

• *Digital Equipment Corporation:* The Distributed Multimedia Group in Karlsruhe, Germany is adapting BBN's public domain ST-II for BERKOM multimedia transport, and investigating ST-II and other protocols for non-BERKOM multimedia projects. They are also sponsoring related research at the University of Karlsruhe. Their platform is an Alpha workstation running OSF.

• *Hewlett-Packard:* Steve Pink has taken a short leave of absence from SICS to act as design consultant for this version. It is currently intended to be a port of the SICS implementation.

• *Siemens-Nixdorf:* We have not received first-hand information about this implementation but understand it to be a part of the BERKOM project.

• *BBN for the IRIS Radio Network:* IRIS will be a mobile telephony system developed for the Canadian Department of National Defense that integrates voice (ST-II) and data (IP) communication in a 600-node internet. IRIS is significant from an ST-II standpoint because each internet gateway must be capable of carrying up to 500 streams simultaneously and rerouting them automatically in the event of node or link failure. Such scales have not been seen in earlier ST-II implementations. This implementation takes BBN's T/20 software as its starting point though it will run on a new platform developed by Computing Devices Canada, Ltd. Special-purpose hardware will route ST-II data packets very quickly based on ST-II's "Hop IDs" embedded in each data packet header. (Hop IDs are analogous to the virtual circuit identifiers, VCIs, defined for ATM.)

The next four implementations have a more research-oriented flavor; each is based on one of the existing public domain versions:

• *Ecole Nationale Supérieure des Télécommunications:* Analyzing the use of RTP over ST-II over ATM for transporting media for existing multimedia applications.

• *FOKUS Research Center Berlin:* Multimedia transport with XTP Layer 4 [15] over the SICS ST-II, aimed at running over ATM. Currently using SunOS with plans to move to Solaris.

• *University of Stuttgart:* ST-II on Sun platforms with FDDI below and XTP Layer 4 above. Port to Solaris operating system likely in the near future.

• *Technical University of Berlin:* Multimedia transport with XTP Layer 4 over ST-II.

**Lessons**

The next section gives some insights from practical experience with ST-II.

• *ST-II Works:* First and foremost, ST-II works. It was designed as a mechanism to establish resource reservations through a multicast tree and then to efficiently transmit data packets through this tree. It accomplishes this goal.

Its way of propagating QoS requests works and its forwarding mechanism is a practical way to transmit multimedia data streams. Furthermore, ST-II requires only unicast routing in the underlying network, over which it builds its own multicast routing. Hence it can be installed on network technologies that do not support hardware multicast or run a multicast routing protocol.

• *QoS Guarantees are Good:* The past year's audio and video broadcasts over the Internet have experienced occasional loss of audio and video signals; such problems are unavoidable over loaded networks without QoS reservation and enforcement mechanisms. We can report from experience with DARTNET and the DSI that networks that support "halfway" QoS schemes such as priority queueing deliver good media quality—until demand exceeds available resources. At that point they fail. We can further report that mechanisms that *enforce* QoS reservations really can and do deliver audio and video streams without those annoying degradations in quality even when demand exceeds capacity.

• *Bandwidth and Delay Guarantees Are Good Enough for Audio and Video*: Current QoS reservations are based on, at most, bandwidth and delay FlowSpecs. These are very crude measures, and there is ample room for improving FlowSpec parameters. But for existing audio and video applications, these crude measures suffice.

• *A "Multimedia" Subset of ST-II has Emerged:* No one has implemented the full ST-II protocol but the implemented subsets are remarkably similar. It would take rather little work to find a baseline implementation that everyone could adhere to, and which would allow a quick transition to interoperability. This is primarily a matter of pruning away the mechanisms designed for research and concentrating on a minimal subset useful for multimedia applications. In addition a few state transitions that the RFC left vague need clarification, and various issues such as the default HELLO packet timing interval need definition. The Heidelberg group has already contributed a great deal of excellent insight in their articles.

• *The RFC Needs Revision:* There is general agreement that the RFC could use substantial revision. It fails to clearly communicate ST-II's basic goals and does not provide motivation for the many mechanisms which it describes. Furthermore, there is a feeling that ST-II is too hard to implement. To a large degree this is due to the RFC not containing a formal definition of states and state transitions. If such a definition were added to the RFC, implementation complexity and time could be greatly reduced. This definition would also increase the interoperability of implementations, since there would be less opportunity for different interpretations of the protocol. And most importantly, a well-defined subset is now ready to move from research into production use for multimedia transport, as noted above.

• *FlowSpecs Should Be Standardized:* Now is the time to standardize on a simple Flow Specification. Although this is still an area of very active research within the Internet community, a simple FlowSpec must be nailed down. Implementors seem in broad agreement that a FlowSpec based on bandwidth and delay, and possibly error characteristics, would suffice for today's audio and video needs. ST-II allows private (non-standard) FlowSpecs which can be used for those applications that cannot use a simpler standard FlowSpec—but there can be no interoperability without a common FlowSpec. Here RFC 1363, "A Proposed Flow Specification," [16] might serve as a suitable guide.

*continued on next page* 

## ST-II Implementations *(continued)*

• *Scalability is Still a Research Problem:* Scalability issues arise in several ways. Each router along a branch of the multicast tree must participate in control packet interactions whenever a target joins or leaves that branch. This can be troublesome if such events happen more often than a router can handle. Routers may have trouble with large numbers of streams, or streams with large numbers of targets, since they must remember state for each stream target in order to perform automatic failure recovery. Finally, a stream's source can also have problems if the stream contains large numbers of targets, or if its target list varies quickly, since it must participate as each target joins or leaves the stream. Finally, it's difficult for routers to reroute many streams at once in the event of link or node failure. This is true for any QoS reservation scheme but ST-II compounds the matter by requiring a flurry of control packets at a critical moment.

**Summary**

ST-II is indeed implementable, at least in subset, and it seems that all the implementors have decided upon roughly the same subset. Furthermore, ST-II works as a successful internet protocol for carrying realtime multicast streams. Indeed it has worked for significant real world applications such as videoconferencing on the Defense Simulation Internet.

ST-II's primary weaknesses? It's too big a protocol for the current needs, its RFC is confusing, and its FlowSpec is too general to allow easy interoperability

We feel that ST-II has in fact proved to be a suitable protocol for realtime multimedia and could be of real practical importance in the near term—if its implementors define a "baseline" subset with a simple FlowSpec suitable for multimedia transmission. This shouldn't be hard.

**Acknowledgements**

We wish to thank the many people who have shared information and opinions with us. These include Debbie Deutsch, Josh Gahm, Varda Haimo, Mark Lefebvre, Craig Partridge, Josh Seeger, and Karen Seo of BBN; Luca Delgrossi, Ralf Herrtwich, and Frank Hoffmann of IBM Heidelberg; Paul Goransson of Wellfleet; Steve Pink of SICS; Gerd Hoelzin and Burkhard Neidecker-Lutz of Digital Equipment Corporation; Spiros Damaskos and Martin Schmeil of FOKUS; Jost Weinmull of TUB; Werner Sinz of the University of Stuttgart; and Frederic Artru of Telecom Paris. The authors, however, must claim responsibility for any inaccuracies and mistakes expressed in this article, and of course for all opinions.

**References**

[1] C. Topolcic, S. Casner, C. Lynn, P. Park, K. Schroder, "Experimental Internet Stream Protocol, Version 2 (ST-II)," RFC 1190, October 1990.

[2] J. Forgie, "ST, A Proposed Internet Stream Protocol," IEN 119, September 1979.

[3] D. Clark, S. Shenker, L. Zhang, "Supporting Real-Time Applications in an Integrated Services Packet Network: Architecture and Mechanism," Proceedings of ACM SIGCOMM, 1992.

[4] L. Zhang et al, "RSVP: A New Resource ReSerVation Protocol," Preliminary Draft, March 1993.

[5] D. Ferrari et al, "Network Support for Multimedia: A Discussion of the Tenet Approach," Technical Report TR-92-072, Computer Science Division, University of California at Berkeley, November 1992.

[6] C. Partridge and S. Pink, "An Implementation of the Revised Internet Stream Protocol (ST-2)," *Internetworking: Research and Experience,* Volume 3, pp. 27–54.

[7] L. Delgrossi, R. G. Herrtwich and F. O. Hoffmann, "An Implementation of ST-II for the Heidelberg Transport System," IBM European Networking Center.

[8] F. O. Hoffmann and L. Delgrossi, "ST-II for the Heidelberg Transport Systems: Protocol, Design, and Implementation," IBM European Networking Center.

[9] R. G. Herrtwich and L. Delgrossi, "Beyond ST-II: Fulfilling the Requirements of Multimedia Communication," IBM European Networking Center.

[10] M. Steenstrup, "Fair Share for Resource Allocation," Draft, December 1992.

[11] L. Zhang, "VirtualClock: A New Traffic Control Algorithm for Packet-Switched Networks," *ACM Transactions on Computer Systems,* Volume 9, No. 2, May 1991.

[12] B. A. Denny, "A Hybrid Algorithm for Combining Best-Effort and Resource Reservation Service," Information and Telecommunications Sciences Center, Report ITAD-8600-TR-93-168, SRI International, June 1993.

[13] J. Gahm, R. R. Hain, and P. Park, "ST-II on T/20 Platform," BBN internal document.

[14] C. Elliott, "High-Quality Multimedia Conferencing Through a Long-Haul Packet Network," forthcoming in Proceedings of ACM Multimedia 93.

[15] "Xpress Transfer Protocol (XTP) Specification," Revision 3.6.

[16] C. Partridge, "A Proposed Flow Specification," RFC 1363, September 1992.

[17] S. Casner, and S. Deering, "First IETF Internet Audiocast," *ConneXions,* Volume 6, No. 6, June 1992.

[18] R. Guido Herrtwich and Luca Delgrossi, "The ST-II Protocol for Multimedia Communication," *ConneXions,* Volume 8, No. 1, January 1994.

**CHIP ELLIOTT** is a Senior Scientist at BBN Systems and Technologies specializing in production software that transmits audio and video streams through packetized networks.

**CHARLES LYNN** is a Senior Scientist at BBN Systems and Technologies. He received his Ph.D. from MIT in 1978 and has subsequently been a lead player in network protocol development for over a decade. He led the effort to implement TCP/IP in the DEC TOPS-20 and has worked in the multimedia area, most recently working on ST-II.

## *Essay*

# Paradigms Found

**by Ed Levinson, Accurate Information Systems, Inc.**

*Interworking* is replacing *interoperating* as the fundamental network computing paradigm. Interworking can be understood through the change it represents in the way people view their workstations. The workstation is being seen as a portal into a collaborative work environment; it is no longer viewed only as a tool or job aid. In the new view the workstation helps individuals work together, or interwork, and in this role places new demands on the software we use. Understanding the paradigm shift will enable product developers to see the critical issues and to provide productivity enhancing software within the new paradigm. This essay examines the causes and implications of the interworking paradigm shift.

**Interworking**

The interworking paradigm is a natural successor to the interoperating paradigm. The interworking paradigm began during the 1980s with the emergence of the TCP/IP protocol suite for data communications, which enabled the creation of large networks and networks of networks, i.e., internets, of computers with diverse architectures. It became obvious, indeed imperative, that the computers on these networks interoperate—transfer data from one to another easily and quickly. Interoperability enabled users to solve problems that required several machines, choosing in turn the one with the resources most appropriate to the sub-task at hand, and moving the data between that machine and the next, to use the next resource. Likewise, groups of workers exchanged information and worked together on tasks and projects when their separate tools could process each others data. These individual tools may be specific programs, specialized hardware, or raw computational power. To take advantage of interoperable capabilities users must know where the computational and data resources are and how to move the data between them.

**Increased power**

During the same decade that interworking emerged Grosh's Law continued to operate. The speed and power of computers continued to increase at the rate of 100 fold over 10 years without significant cost increases. This trend resulted in the growing use of workstations—powerful computers dedicated to individuals. These workstations took advantage of the network to leverage their computing power with applications and data available from server hosts that were also on the network. The result today is that more workers, either alone or in groups, are doing more computing in more places.

Similarly data communication capacity (bandwidth or speed) has increased 10 fold over the last decade and will continue to do so over the next. These capacity increases are taking place in both wide area networks (WANs) and local area networks (LANs). Capacity differentials will continue to exist, LAN speeds will continue to outpace WANs and will approach the speed of backplanes in older microcomputer architectures. At those speeds it is even more likely that data and computing resources need not be co-located or even be inside a single machine.

The increased computing power is available at prices that make it attractive to place this power on people's desks. Some of the additional computing power is being used to make the machine easier to use, through user interfaces that employ graphics, voice, and video. This new desktop capability will enable ideas to be communicated with greater ease through pictures, charts, and engineering drawings.

To permit this to happen will require the communications capability be matched to the higher data rates that multi-media data require and that new data communications will provide.

**Data availability**

The increased data communication capability makes the information, used by many departments within an enterprise, widely available. This wide availability will, in turn, generate new incentives to make the data consistent across all departments and to promote data sharing between them. People will expect the system take over the tasks of locating the resources and moving the data. Data sharing can then be based on policy parameters, unconstrained by technical barriers.

New applications, distributed in nature, will partition the work across the network of computers and execute different tasks concurrently. They will enable users to work together in ways that cannot yet be predicted. Some impact of increased computer power and network speed can already be seen. For e-mail, increased power has led to increased mail system sophistication. New standards such as MIME and X.400(88) provide for the exchange of non-textual data, e.g., images, voice, binary files, etc. Within these environments, additional work is required to enable people to exchange complex objects without being tied to specific applications. As open networking protocols have created an energetic market for interoperable network hardware and software, open application protocols offer a similar possibility for applications.

Increases in delivery speed will allow workers to interact in near real time; geography will be minimized as a barrier to rapid interaction between physically separated colleagues. More sophisticated distributed applications, seen today in workflow and workgroup applications will quicken the pace of work and will lead to more work being done concurrently. Increasingly, people will find themselves at their desks participating in concurrent collaborative efforts conducted through the workstation and over the network infrastructure that connects them together.

**Paradigm shift**

This, then, is the paradigm shift. A change in the way workstations are used; a change from the workstation as a tool to the workstation as the medium for communication and collaborative work. The workstation's role as a tool or job aid will not become subservient to its new role, it will do both equally well. The workstation's additional power will be used to facilitate the large communication effort required for collaborative work. Producing a paradigm shift from Interoperating to Interworking.

**Acknowledgement**

This essay is one result of many stimulating discussions with Einar Stefferud, Network Management Associates, Inc., and the members of the Information Management Directorate, US Army Armament Research, Engineering, and Development Center, Picatinny Arsenal, NJ. The author gratefully acknowledges their contribution and encouragement. This work was prepared in part under US Army Contract DAAA21-88-D-0025.

**EDWARD LEVINSON** holds a B.S. from Washington University in St. Louis and an M.S.E. and Ph.D. from the University of Pennsylvania. Currently he is at Accurate Information Systems, Inc. where he works on electronic mail issues and infrastructure planning. Prior to joining Accurate he was at Syntrex, Inc., where he had network and software engineering responsibilities. He began his career developing software at Bell Laboratories, Inc. and, at divestiture, joined Bellcore. His current passions are embedding complex structures in MIME and being able to address e-mail to anyone, anywhere, as long as they're connected. He can be reached as: `elevinso@accurate.com`

*Essay*

# Paradigms Lost

**by Einar Stefferud, Network Management Associates, Inc.**

## Introduction

Following my attendance in 1987 at the *First TCP/IP Interoperability Conference* (the first INTEROP) in Monterey where relationships between the Internet Community and the ISO/CCITT Community could be seen to unravel right before our very eyes, I wrote:

> "A Plea For Internet Peace," in *ConneXions,* Volume 1, No. 5, September 1987, pp. 13–15.

My objective then was to foster convergence and interworking among deployed protocols. Then we called it "transition" because we believed that ISO/CCITT protocols held some kind of magic advantage so that convergence would simply mean transition *from* whatever we have *to* OSI. Then I thought the market would soon put a high value on Formal International Agreement on Protocols.

With hindsight I sometimes regret writing my plea because so many efforts since then have focused on compromise of important principles in the interests of peace. For me, the notion of compromise between the two communities has now lost its allure.

## Paradigms, paradigms

After more than 6 years of pondering, I have come to a clear conclusion; Our two communities have been dueling with each other across *paradigm boundaries* and this has not been a rewarding endeavor. In this discussion, a paradigm is a "reality model" in which I express my thoughts in language which reflects the semantics of my reality. Choose a different paradigm, and you will be using different semantics in your language, and be operating in a different reality. Operating in different realities together is generally difficult.

Here is a map of my paradigms to show you what I see:

– *Computing Paradigm: One ended networking.*

- This is the stuff of Application & Operating System developers.
- Of little or no interest to Networkers or Internetworkers.
- Of course, lots on non-networkers are very interested!
- Here, Networking is Somebody Else's Problem (SEP).

– *Networking Paradigm: Two Ended Networking.*

- Interesting to Real Networkers, but not to Internetworkers.
- All Interesting Network Connections have at least Two Ends.
- Assumes a Single Owner for all components in the network.
- Connectivity and Control are Key Motivators.

– *Internetworking Paradigm: Two Autonomously Owned Ends.*

- Very Interesting to Internauts but often not to Networkers.
- Internets involve Distributed Ownership and Control.
- The other end simply cannot be trusted! Owned by Someone Else!
- Failure must be Expected and Planned for.
- Interoperability is a Key Motivator.

– *Interworking Paradigm: Unbounded Autonomously Owned Ends.*
- Interesting to Application Users.
- Above the Application Layer—Goes beyond Interoperation.
- Requires Creation, Transfer, and Usefulness of Received Application Information Objects.

OK—So what does this paradigm map tell us? Before drawing any conclusions let's see who is living in which paradigms. Then maybe we can see why our conversations are so bizarre.

**Computing**

LAN Operating Systems folks live in the *Computing Paradigm.* Basically, they have subsumed the LAN "network" into the "backplane" and this solves networking problems by making them go away; those are somebody else's problems. LAN folk spend lots of time talking about APIs as though APIs solve communication network problems.

Internauts spend little time on APIs. Application Programming Interfaces primarily have to do with Application program portability, not data communication.

**Networking**

Networkers (SNA and DECnet and X.25 and ISO/CCITT folk) live in the *Network Paradigm.* They assume that they own it all, whatever the reasoning, or that what they don't own is out of scope. If it isn't in their territory, it must be somebody else's problem.

Don't leap to conclusions here without stopping to think about it. Ponder a bit. "Owned" can take on different meanings depending on whether you are operating a network, developing product lines, or writing protocol specifications.

ISO and CCITT behave as though they own the entire intellectual space inside the scope of their specifications, and they avoid all reference to anything outside that space. The word "gateway" is not in their lexicon. Gateways are things that live outside their scope; gateways engage in protocol translation between their problem space, and somebody else's problem space. When we make something into someone else's problem, it disappears.

**Internetworking**

Internauts live in the *Internetworking Paradigm;* knowing, designing, and building networks of networks where one must always assume the other end of any connection/association will be owned and controlled by someone else. The Internaut knows that there is no hope of getting by with a declaration that something at the other end is "out of scope" or that "it is not in our lexicon." Mutual collaboration is the only useful tool in an Internet.

In the Internet, making something into somebody else's problem does not simply make it disappear. An Internaut might say "We have met somebody else, and they are us!"

**Interworking**

And, pity the poor End User, who lives in the *Interworking Paradigm,* and only wants to get some work done by usefully exchanging Application Information Objects with a multitude of other End Users who work on a multitude of systems and networks whose design and administration they cannot control, or even influence.

**Mixing paradigms**

Now, imagine getting randomly selected people together, each living in one of these different paradigms, to draw up a plan for the future! What language should they try to speak? What common reality model should they try to use to communicate? How can they communicate?

*continued on next page* 

## Paradigms Lost *(continued)*

Let me provide a concrete example, from the land of network management. In ISO/CCITT terms, "integrated" management typically means organizing things to get all the elements under control of some Lowest Common Manager. The "Containment Tree" is a central concept in X.700 (also known as CMIP/CMIS), and it pretty much models the world as an authority tree.

In the Internet "integrated" means finding ways to manage without actually getting everything together under some Lowest Common Manager. Any Internet solution requires dealing with distributed authority and distributed responsibility. Internauts cannot solve an autonomous manager's problems by making autonomy go away, or by declaring some arbitrary line of demarcation.

No one is in control of new connections to an Internet, and the Quality Of Service (QOS) thus depends on both ends of any connection pair being well managed!

Note that the same word (integrated) now seems to have opposite meanings in our two paradigms (Networking and Internetworking). No wonder the conversation is so strained and unpleasant. It is interesting how this difference affects fundamental engineering design principles in our different paradigms.

### Impact on engineering

Internauts learned (and continue learning) that it is not a good idea to trust anyone outside your domain of control, which in the Internet is amazingly small. Internauts expect things to fail, and plan for it. Internauts don't trust their networks and their protocols don't trust the other end.

Networkers tend to trust their networks with deliberate intent. They put great importance on network trustworthiness. In some cases, for example, they have tried to provide host security by securing the entire network. This works in severely limited enclaves, but not in an unbounded Internet where there is no way to decide who to exclude from interconnection.

Besides, we all know that the many security threats come from within. How do you suppose IBM might secure its corporate SNA network from its own employees?

In the X.700 case, the management protocols require fully Reliable Transport Services to carry all Protocol Data Units. In SNMP, the protocol uses User Datagram Services to carry Protocol Data Units. On balance, which of these do you think will fail first in a misbehaving network?

And, in the Interworking Paradigm, End-Users can't trust anyone! More than likely, the e-mail recipient of an Application Information Object will find that it cannot be processed for one of many possible reasons.

### The paradigm shifting process

Shifting between paradigms is an "AHA!" kind of experience. At least some who read this message will be in the middle of an AHA as they read it. It often happens with a POP. Or perhaps a little slower, as with the Star Trek Transporter, but there is no mid-point where you can stop to camp for the night. The mind seems to work best when you are in one paradigm or another, but not stuck somewhere in between, or flipping between them.

Schizophrenia anyone?

I sometimes regret my "Plea for Internet Peace," because it may have contributed to Internet efforts to compromise with OSI proponents when there is no basis for compromise across paradigm boundaries.

Pick any pair of paradigms; perhaps you can find some in another universe. Compromising between paradigms is pretty much like a trip halfway through the Start Trek Transporter. [It's not a pretty sight!] This observation has nothing specific to do with TCP/IP or ISO/CCITT. Consider religious arguments, which tend to have the same "Look and Feel."

What should we do to resolve our various paradigm differences? A first step is to realize where we are standing, relative to each other, and realize the nature of what we are committed to accomplishing—namely trying to get along in different paradigms together, which is very much like getting along in different religions together.

Paradigm shifting is like religious insight. The AHA is often accompanied by strong emotional reactions when we suddenly find ourselves with whole new views of the universe. A good description of it is "a period of luminosity."

An interesting social result is that those who have popped tend to suddenly talk in a strange new language, and start trying to convert their old friends and colleagues to the new paradigm, but the old friends "just don't get it!" Let's face it—How many of you out there are convinced (even now) that I must be crazy? One aspect is that us popped people seem to adopt a superior attitude. We are, after all, able to see things much more clearly now. Our main problem is that unpopped people "just don't get it!"

Over the years, I have read many messages related to our TCP/ISO wars. Having worked with people in all four paradigm camps over the last 32 years, I see that what we need most is to recognize that we are driven by very different commitments about what is important. That we use the same words in different meanings, and generally float about on different planes together, trying to hold hands.

**A recommendation**

My recommendation is that those who want the *Internet* to adopt their pet *Networking* technologies, should come on over, and get involved in trying to solve the prevailing *Internetworking* problems by working shoulder to shoulder in the IETF Working Groups. Sooner or later they will pop into the Internet Paradigm, and live happily ever after...

Of course, the reverse holds just as well, though I admit it is really hard to work in a paradigm that one no longer believes is interesting. That is one of the really serious problems with *Popping* into a new paradigm. It is really hard to go back. One more thought: Paradigm Popping is only done by individuals, not by groups.

**Conclusions**

I should carefully note that discovery of new, perhaps higher level paradigms does not invalidate the old ones. It is more a case of suddenly becoming aware of a larger reality, in which old paradigms remain valid, but are now parts of a larger whole.

Computing happens inside networks which make up internets, which in turn will provide the infrastructure for our working together. Working toward Convergence is a long, slow process of drawing people into some new paradigm, where the incentives are to find singular common solutions that meet the requirements of an open marketplace for open interconnection, open program portability, and open interworking environments.

*continued on next page*

## Paradigms Lost *(continued)*

Our best operational example of this marketplace to date is our beloved Internet, with its absence of central control, and which is not susceptible to anyone cornering our Electronic Real Estate.

The Internet paradigm may be the result of an accidental collision of academic research and military needs, but its deeper meaning is that it defines new possibilities for open interworking on a different plane, without singular central control, or a single point of failure. It represents a whole new world of possibilities for interworking, and its requirements will define where we eventually converge.

In this context, formal international political agreements cannot mean very much, regardless of how selfless and well-meaning are the people who labor to generate the specifications. The "Internet as a Marketplace" will eventually make the decision for us.

**References**

[1] Stefferud, E., "A Plea For Internet Peace," *ConneXions,* Volume 1, No. 5, September 1987, pp. 13–15.

[2] Jacobsen, Ole, "The Trouble with OSI," *ConneXions,* Volume 6, No. 5, May 1992, p 62.

[3] desJardins, Richard, "*Opinion:* OSI is (Still) a Good Idea," *ConneXions,* Volume 6, No. 6, June 1992, p 33.

[4] Metcalfe, Bob, Smart, Bob, and Blackshaw, Bob, "Letters to the Editor," *ConneXions,* Volume 6, No. 6, June 1992, p 37.

[5] Rose, Marshall, "Comments on: '*Opinion:* OSI is (Still) a Good Idea,'" *ConneXions,* Volume 6, No. 8, August 1992, p 20.

[6] desJardins, Richard, "Comments on: 'Comments on: *Opinion:* OSI is (Still) a Good Idea,'" *ConneXions,* Volume 6, No. 10, October 1992, p 43.

[7] desJardins, Richard, "Internet 2000," *ConneXions,* Volume 6, No. 10, October 1992, p 24.

[8] Hoffmann, Harald, "Letter to the Editor," *ConneXions,* Volume 6, No. 11, November 1992, p 31.

[9] Hoffmann, Harald, "Letter to the Editor," *ConneXions,* Volume 7, No. 10, October 1993, p 23.

[10] Rose, M.T., "Network Management: Status and Challenges," *ConneXions,* Volume 7, No. 6, June 1993, p 11.

[11] Crocker, David, "Letter to the Editor," *ConneXions,* Volume 7, No. 2, February 1993, p 24.

**EINAR STEFFERUD** ("Stef") is Founder (1969) and President of Network Management Associates in Huntington Beach, California. His firm provides strategic technical advisory information services to internetworking clients ranging from universities to aerospace, army laboratories, network service providers, and government agencies. He has long been active in the IETF and other professional activities. He serves as Chair of IFIP Working Group 6.5 and is Chair of the EMail World Conference. He was listed by *Communications Week* among their Top 10 Visionaries for 1993. Stef used to say that you could send him electronic mail to `stef@<any host name on the ARPANET>`, but that was in the Good Old Days. Today, you can reach him most reliably as: `stef@nma.com`

## Letter to the Editor

Editor, *ConneXions,*

I read with interest the recent review of one of my books: Quarterman, John S. and Susanne Wilhelm, *UNIX, POSIX, and Open Systems: The Open Standards Puzzle,* Addison-Wesley, 1993. [Ed.: See *ConneXions,* Volume 7, No. 10, October 1993, page 24 for the review].

I take no issue with the comments about the book itself; everyone is entitled to their opinion.

I would, however, like to respond to the interspersed editorial commentary about standards. I am enlightened to learn that Microsoft is a good model of standards making. That must explain why the MS-DOS command I remember best is "control-ALT-delete," and why Windows-NT appears to ignore every lesson learned from UNIX.

On the other hand, traditional standards processes can easily become waterlogged with streams of process until they sink or simply lose the race. OSI was going to replace TCP/IP in the next two years for a decade, and hasn't yet.

However, there is a third way. Computer networks permit making standards in a relatively informal manner, emphasizing implementation during specification and before standardization, and with wide participation. Perhaps the best example of this is the IAB or IETF standardization process, which has led to an exponentially growing internetwork based on the most widely implemented set of network protocol standards in the world.

Incidentally, readers of *ConneXions* might be interested to know that the longest chapter in the book is about the IAB standards process.

Thanks,

—*John S. Quarterman, Texas Internet Consulting*
`jsq@tic.com`

## Write to *ConneXions*!

Have a question about your subscription? Are you moving, and need to give us your new address? Suggestions for topics? Want to write an article? A letter to the Editor? Have a question for an author? Need a *ConneXions* binder? Want to enquire about back issues? (there are now over 80 to choose from; ask for our free 1987–1992 index booklet and the 1993 index sheet). We want to hear from you. Send your questions, comments or suggestions to our new address:

*ConneXions—The Interoperability Report*
303 Vintage Park Drive
Suite 201
Foster City, CA 94404–1138
USA
Phone: +1 415-578-6900 or 1-800-INTEROP (Toll-free in the USA)
Fax: +1 415-525-0194
E-mail: `connexions@interop.com`

## Book Review

*Distributed Computing: Implementation and Management Strategies*, R. Khanna, Editor, ISBN 0-13-220138-0, Prentice-Hall, 1994.

**Another palliative**

Every few years, our fast-moving profession invents a term to describe the next wave of technology to solve all our problems. Things never quite live up to the promise, of course, but that does not stop the tidal wave of media enthusiasm. Such has been the fate of "Distributed Computing." Once you get past a trivial definition, it is never quite clear what is *excluded* by the term, what is accomplished by the technology, or how to go about getting it and using it. This book tries to remedy these deficiencies. Given that the topic is not subject to quick or clever description, the book succeeds quite well. It is an exposition of the ideas, technologies and techniques that can be banded together into an enterprise-wide, computational service, and it is rich with discussion about approaches toward developing such a service.

The most common terms for the technologies described in the book are "middleware" and "enterprise networking." They attempt to enhance raw network transport services by providing additional building-blocks for development of distributed applications and for operation of networked platforms. Broadly, these encompass three nearly independent areas of service: configuration and query, networked operating system, and application development. The first allows a member of the distributed environment to obtain information about itself and other resources on the net. The second allows multiple computers to cooperate to form an integrated service, e.g., distributed time and remote file access. The last facilitates development of end-user applications which are, themselves, distributed among various networked systems. With such a range of tasks to perform you might expect "distributed computing" technology to be complicated to deal with. You'd be right.

**Organization and style**

*Distributed Computing: Implementation and Management Strategies* is a series of independently-written expository chapters. Each is serious and relatively information-dense. Don't look for clever asides or outrageous myth-busting diatribes. Many of the chapters are written by principals in the development or promotion of the respective technologies, so skeptical analysis is not generally present. This does not mean that reality is absent. In a chapter on "Migration Strategies," Bob Morgan cautions: "DCE's complexity is a challenge in itself, however. It is likely that only organizations that are already hip-deep in the distributed computing swamp will be able to appreciate the benefits..."

After an introductory chapter, the book is divided into 3 parts and a series of Appendices which really are short chapters forming a bonus section. In Part 1 six chapters discuss basic technologies, including the much-heralded DCE, as well as the ONC+ enhancements to the much-used ONC/NFS. Jeff Schiller's chapter on Distributed System Security is as good an introduction to the topic of network-based security as I've seen. Security discussions tend to sound like black-magic incantations to the non-expert, yet Jeff covers the ground thoroughly and clearly. (And you probably won't miss much if you skip the few math notations.)

The second Part of the book comprises 5 case studies, in keeping with the book's effort to be pragmatic. Included are descriptions of CMU's *Andrew* and MIT's *Athena* projects, as well as discussion of "institutional" implementation and use projects at the University of Michigan, Hewlett-Packard, and Eastman Kodak. All of these chapters are useful; the Andrew and Athena ones are particularly excellent for the perspectives they offer.

Part 3 discusses "strategies" and "issues." In other words, these chapters try to step back and discuss things using a broad perspective. The first of these is a downright clever chapter in which Khanna got vendor strategy statements from Sun, HP, IBM and Microsoft (with one of the appendices coming from Apple). Page-flipping among these consecutive summaries was particularly interesting for comparing their architectural charts. If you have any expectation that this realm of technology will make your life simpler, the charts should finish that off.

**No quick answers**

Media discussion about distributed computing has heavily emphasized OSF's work on DCE and DME. Both are covered in detail in the book, with Morgan's chapter discussing adoption and use. As long as you remember that these real experiences were by highly skilled staff in a highly advanced environment, the discussion adds flesh to the otherwise-spare repertoire of knowledge about DCE realities.

Distributed computing technology provides an enhanced infrastructure, rather than direct solution to specific end-user requirements. This books does the same thing. It provides extensive detail *about* an area of new technology and service and about the approaches that can be taken to incorporate the capabilities. It does not contain any cookbooks or procedural checklists. It provides a basis for doing extended analysis and a discussion of the current industry experiences. Like the move to distributed computing, reading the book will be an investment in the future. If you develop distributed systems or must make acquisition decisions for them, I recommend the investment.

*—Dave Crocker, Silicon Graphics*

## NetWorld®+Interop® 94 World Tour

Mark your calendar for the following five dates and locations for NetWorld+Interop 94:

| | |
|---|---|
| *NetWorld+Interop 94, Las Vegas, Nevada:* | *May 2–6, 1994* |
| *NetWorld+Interop 94, Berlin, Germany:* | *June 6–10, 1994* |
| *NetWorld+Interop 94, Tokyo, Japan:* | *July 25–29, 1994* |
| *NetWorld+Interop 94, Atlanta, Georgia:* | *September 12–16, 1994* |
| *NetWorld+Interop 94, Paris, France:* | *October 24–28, 1994* |

NETWORLD®+INTEROP® 94

## The First Annual "Ima" Awards

The Internet Multicasting Service is pleased to announce the results of the highly prestigious and somewhat obscure "Ima" awards for 1993. The awards recognize people, software, and computers that have made a significant contribution to the Internet. In these days of Internet hype, our hope is to call attention to a few of the people making truly significant contributions.

The "Ima" awards are selected by our Prize Committee, an anonymous group of distinguished users and developers of the Internet. The "Ima" awards will be given once a year. The award classifications are not fixed and the Prize Committee will decide each year which categories are appropriate.

**Lifetime Contribution to Basic Research**

The "Ima" Award is presented to Dr. David Clark in the category of "Lifetime Contribution to Basic Research." Dr. Clark has a long and distinguished career in basic research and has made significant contributions in the areas of computer security, protocol design and implementation, and high-speed networks. Most recently, Dr. Clark has been a significant contributor in research in the areas of service classes, resource reservation, and policy routing in the Internet. Perhaps the most significant contributions from Dr. Clark have been in the field of research known as "herding elephants," showing us how to avoid giant bureaucracies and keep the Internet working on rough consensus and running code.

**Geek of the Year**

The "Ima" Award is presented to Dr. Stephen Deering in the category of "Geek of the Year." Dr. Deering has been instrumental in the fields of mobile computing, multicasting, and has been a leader in development and deployment of a next generation Internet Protocol. This leadership across a wide variety of research fields has been a significant contribution to the development of the Internet.

**Most Innovative NIC**

The "Ima" Award is presented to the RIPE NCC in the category of "Most Innovative NIC." RIPE has proven to be a truly innovative leader, showing an unparalleled flexibility in assigning address space to country-level NICs and in actively establishing new NICs around Europe. The RIPE NCC has actively supported the quarterly RIPE conferences in Europe and has shown global technical leadership in the deployment of *Classless Inter-Domain Routing* (CIDR), BGP-4, and other crucial technologies for the Internet.

**Most Innovative Regional Network**

The "Ima" Award is presented to the *Widely Integrated Distributed Environment* (WIDE) Project (Japan) in the category of "Most Innovative Regional Network." The WIDE Project has consistently shown an ability to involve a wide population of users and researchers in the Japanese Internet and has shown global technical leadership in the areas of deployment and implementation of ISDN, for support of TCP/IP on a wide number of unusual platforms, in mobile computing, and in support for character sets.

**Most Innovative Application**

The "Ima" Award is presented to NCSA *Mosaic* in the category of "Most Innovative Application." 1993 has been the year of the Internet Application, but the National Center for Supercomputer Applications has risen above this distinguished fray to develop a truly impressive piece of programming. NCSA *Mosaic,* which runs on X Windows, the Macintosh, and Microsoft Windows, is an interface to the *WorldWide-Web,* but has also integrated transparent access to other Internet services, ranging from FTP to WAIS to *Gopher.* NCSA *Mosaic,* in addition to an unparalleled flexibility in design, has proven to be a superb implementation.

**Most Innovative Architecture**

The "Ima" Award is presented to Tim Berners-Lee and the *World-WideWeb* (WWW) Team in the category of "Most Innovative Architecture." By carefully studying the areas of multimedia, hypertext, and SGML, WWW has developed a practical and powerful paradigm for finding and deploying information on the global Internet. The explosive growth of data on the Internet in 1993 is a tribute to the architectural vision embodied in the World Wide Web.

**Most Exhaustive Documentation**

The "Ima" Award is presented to the O'Reilly & Associates' "Nutshell Series" in the category of "Most Significant Contribution to Internet Documentation." While there have been many books and papers on the Internet, the Nutshell Series has proved to be an in-depth source of technical information through books dealing with such arcane topics as *Managing NFS and NIS, DNS and BIND,* and *Sendmail.* The Internet has gone far beyond the stage of oral history and high-quality documentation from sources such as the Nutshell Series has become an essential part of our infrastructure.

**Most Effective Bureaucrat**

The "Ima" Award is presented to Dr. Stephen Wolff in the category of "Most Effective Bureaucrat." With the advent of the HPCC, NREN, NII, and other programs, the Internet has become a highly visible public policy area. Despite intense public visibility and an often contentious policy environment, NSF has continued to take a leadership role in the Internet, funding and managing programs in a large number of key infrastructural and research areas. It is a tribute to Dr. Wolff's leadership that NSF has maintained such a visible role in the Internet.

**Honorable Mentions**

Honorable Mention "Ima" awards go to *WIRED* for the coolest new magazine, Bo Pitsker of the INTEROPnet as "network operator of the year" for his ability to swallow a firehose and deliver a several-thousand node network three times per year, and the *On-Line Bookstore* (OBI) as the "most innovative new service" for bringing people like Stephen King onto the Internet.

## To our readers

Because of company reorganization and relocation, you may not have received a *ConneXions* subscription renewal letter. Please take a moment to check the expiration date for your subscription. This is printed on the mailing label. If your expiration date is soon, please mail in your payment or send us e-mail with the appropriate information. You can also help us greatly by including your payment with your subscription renewal or new order. We often get separate payments without further identification and this leads to lots of hard detective work. If your accounting department is mailing us payment separately, please ask them to mark the check with some kind of reference information ("Subscription renewal for Joe Smith"). If you are taking over a subscription from another individual at your company, please indicate this. Please mark any name or address changes clearly. We are developing a new subscription database which should be coming on line in the next couple of months. While we can't offer sweepstakes or free clock radios, the new system should make the subscription process a lot more efficient. Thank you for your patience.

*—Ole J. Jacobsen, Editor and Publisher*

CONNEXIONS
303 Vintage Park Drive
Suite 201
Foster City, CA 94404-1138
Phone: 415-578-6900
FAX: 415-525-0194

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

ADDRESS CORRECTION
REQUESTED